Final Report 2000-07

PB2000-107334

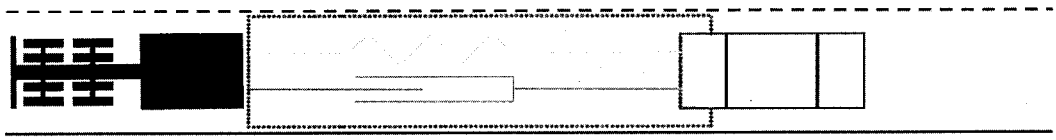

Radar Based Longitudinal Virtual Bumper Collision Avoidance System Implemented on a Truck

Minnesota Guidestar

Technical Report Documentation Page

1. Report No.	2.	3. Recipients Accession No.
MN/RC - 2000-07		

4. Title and Subtitle	5. Report Date
RADAR BASED LONGITUDINAL VIRTUAL BUMPER COLLISION AVOIDANCE SYSTEM IMPLEMENTED ON A TRUCK	March 1999
	6.

7. Author(s)	8. Performing Organization Report No.
Alec Gorjestani Max Donath Lee Alexander	

9. Performing Organization Name and Address	10. Project/Task/Work Unit No.
University of Minnesota Department of Mechanical Engineering 111 Church Street SE Minneapolis, MN 55455	11. Contract (C) or Grant (G) No. c) 74708 wo) 59

12. Sponsoring Organization Name and Address	13. Type of Report and Period Covered
Minnesota Department of Transportation 395 John Ireland Boulevard Mail Stop 330 St. Paul, Minnesota 55155	Final Report
	14. Sponsoring Agency Code

15. Supplementary Notes

16. Abstract (Limit: 200 words)

This report describes results from a series of experiments using the virtual bumper collision avoidance algorithm implemented on a Navistar tractor cab.

The virtual bumper combines longitudinal and lateral collision avoidance capabilities to control a vehicle in normal and emergency situations. A programmable boundary, the virtual bumper, defines a personal space around the host vehicle. Researchers used a radar and a laser range sensor to sense the location of vehicles in front of the truck. Target vehicles that enter the truck's personal space impose a virtual "force" on the host, which in turn modifies the vehicle's trajectory to avoid collisions with objects in the field of view.

Researchers tested the virtual bumper longitudinal controller under different driving situations and at different speeds. The experiments included several scenarios: Adaptive Cruise Control, the truck performing a critical stop for a stationary target vehicle, and situations that simulate stop-and-go traffic.

Results from the virtual bumper longitudinal experiments were favorable. The algorithm demonstrated robustness to sensor noise and the ability to maintain a safe headway for both normal and emergency driving scenarios. Researchers currently are improving the sensing technology and incorporating a road database, which contains roadside features to greatly reduce, if not eliminate, false target detection.

17. Document Analysis/Descriptors		18. Availability Statement
Collision avoidance	Headway control	No restrictions. Document available from:
Crash avoidance	Heavy vehicle longitudinal control	National Technical Information Services,
Longitudinal control		Springfield, Virginia 22161

19. Security Class (this report)	20. Security Class (this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	107	

RADAR BASED LONGITUDINAL VIRTUAL BUMPER COLLISION AVOIDANCE SYSTEM IMPLEMENTED ON A TRUCK

Final Report

Prepared by:

Alec Gorjestani, Lee Alexander and Max Donath

Department of Mechanical Engineering
University of Minnesota
111 Church Street SE
Minneapolis, MN 55455

March 1999

Published by:

Minnesota Department of Transportation
Office of Research & Strategic Services
Mail Stop 330
395 John Ireland Boulevard
St. Paul. MN 55155

This report represents the results of research conducted by the authors and does not necessarily represent the views or policies of the Minnesota Department of Transportation. This report does not contain a standard or specified technique.

Acknowledgments

It would be difficult to complete such an involved project without the generous help of other people. We would like to take this opportunity to thank them for their generosity.

Craig Schankwitz, Sameer Pardhy, Bryan Newstrom, and HeonMin Lim deserve thanks for driving the target vehicle during the experiments and their flexibility in allowing us to schedule considerable time aboard the Navistar tractor. They were a valuable resource for technical information. We also wish to thank Ted Morris for his assistance and prompt responses to numerous queries about his PIC software and VxSim.

We would also like to thank Jack Herndon and others at Mn/ROAD for their flexibility and assistance with the experiments. Jack was very helpful in scheduling drivers and coordinating the use of the mainline highway for the higher speed experiments.

We wish to acknowledge the ITS Institute, the Center for Transportation Studies, Mn/DOT, and the U.S. DOT for their partial support of the effort described here. We also need to thank the Navistar International Transportation Company for providing the International Navistar 9400 truck that was used for the experiments.

Table of Contents

LIST OF FIGURES

Executive Summary

In this report, we describe the implementation of the virtual bumper collision avoidance algorithm for highway vehicles. The virtual bumper combines longitudinal and lateral collision avoidance capabilities to control a vehicle in normal and emergency situations. A programmable boundary defines a "personal space" around the "host" vehicle. Incursions by "target" vehicles or other objects into this space are sensed by range sensors mounted on the host vehicle. This virtual deflection of the bumper defined by the boundary generates a virtual "force" on the host which affects its trajectory in a manner that attempts to avoid (or at least mitigate) collisions. The relationship between the virtual bumper deflection and the virtual force that is applied to the host vehicle is computed based on a non-linear relationship which is a function of the range and the derivative of range to the objects in the host vehicle's environment. The road also induces a virtual force, which attempts to keep the host within its lane of travel.

The virtual bumper consists of three main subsystems. The longitudinal control subsystem incorporates impedance control to adjust the headway to proceeding vehicles and maintain the desired traveling speed when no obstacles are present. The lateral control subsystem is an impedance controller that maintains the host vehicle's position in the lane and performs collision avoidance in the side regions of the vehicle. The final component of the virtual bumper is the lane change subsystem, which determines the safest lane of travel based on the road environment and issues lateral position commands that perform lane changes (or direct the vehicle off the road if needed). Host vehicle velocity and acceleration is measured using a Differential Global Position System (DGPS) and then used in the collision avoidance controllers. In this report, we describe the implementation of the longitudinal controller subsystem and demonstrate its ability to operate under several different types of scenarios. These include (a) Adaptive Cruise Control (ACC) on the highway, (b) a major traffic slowdown, (c) a critical stop situation, and (d) stop and go traffic.

The virtual bumper longitudinal controller was implemented first in simulation and then on an International Navistar 9400 tractor and tested under several driving situations and at several speeds. Simulation tests were first used to demonstrate that the algorithm is robust to range sensor noise and DGPS noise. DGPS noise did significantly raise the virtual forces. Even though this was not a serious issue, future DGPS used for the system will have significant less noise, i.e. higher accuracies, lower variance, and faster sampling rates.

Experiments on the truck were performed at the Minnesota Road Research Facility (Mn/ROAD). We tested the ACC functionality of the controller by setting the host cruising speed at a higher speed than that of the target vehicle. After the range sensors (radar and lidar) detected the target vehicle, the virtual bumper successfully slowed the host vehicle to match the speed of the proceeding target vehicle while maintaining the desired headway.

A major traffic slow down scenario was used to determine the controller's ability to comfortably decelerate the host when the target was traveling at a significantly slower speed. The longitudinal controller was able to slow the host vehicle (using brakes) and establish the desired headway with reasonable deceleration levels.

To simulate a critical stop scenario, the target vehicle was parked in the host's lane of travel. The host was then driven toward the target at varying speed. The controller successfully stopped the host at a safe headway, but we discovered that the brake actuation did not have the required sensitivity to apply the brakes gently at low velocities. This caused the truck to stop abruptly at the final stage of the maneuver instead of gently approaching a full stop. This result was due to limited access to the brake controller imposed by the brake control supplier.

The virtual bumper was also tested in scenarios in which the target suddenly slowed down, speeded up or was simulating stop and go traffic. The system was able to maintain a safe headway for all of these driving scenarios. The aforementioned brake actuation hardware

limitations did affect the stop and go experiment in the same manner as the critical stop experiments.

The algorithm demonstrated robustness to sensor noise and the ability to maintain a safe headway for both normal and emergency driving scenarios. While we were pleased with the overall performance of the longitudinal controller, we did notice that the algorithm was sensitive to false targets produced by the range sensors and velocity tracking errors of the velocity controller due to hardware limitations. Continually improving sensing technology and road databases which contain roadside features can be used to greatly reduce, if not eliminate, false target detection. We plan on improving the brake actuation by gaining direct access to the brakes. This should greatly improve the velocity controller's ability to track velocity commands.

The positive result of the virtual bumper longitudinal experiments demonstrated that we are able to successfully integrate sensing and control of steering, throttle, and braking as part of a collision avoidance system, at least from longitudinal situations affecting the front end of the vehicle. We are ready to move on to the next phase of the development of the virtual bumper. Future work includes: (a) developing methods for integration of multiple radar units and other similar sensor inputs to inform us about what exactly is around the host vehicle, (b) integration of the virtual bumper with knowledge of the road environment to include what is on and adjacent to the road and what each lane of the road is designed for, (c) continued development and testing of a virtual bumper for lateral control of a vehicle, (d) consideration of rear directional sensing and suitable algorithms to handle vehicles bearing down on the host from the rear, (e) integration of the longitudinal and lateral virtual bumper subsystems, and (f) developing appropriate human interfaces to the virtual bumper.

Chapter 1

Introduction to the Virtual Bumper

1. Introduction to the Virtual Bumper

The virtual bumper is a collision avoidance strategy that combines longitudinal and lateral collision avoidance capabilities to control a vehicle in normal and emergency situations. The strategy has been successfully tested in simulation and has demonstrated the ability to avoid collisions in a variety of situations [1][2].

This thesis documents the implementation of one subsystem of the virtual bumper collision avoidance scheme (the longitudinal controller) on an International Navistar 9400 tractor. First, a simulation is used to provide insight into what effects various parameters and real world issues have on the algorithm. Second, data collected from a variety of experiments is presented and compared with the simulation.

The motive behind this thesis is to design, test and evaluate the longitudinal controller on a real vehicle. We worked with commercially available hardware to validate the theory tested in simulation and to verify that the desired response is achievable on an actual truck. Sensors are continuously improving. In this implementation we used radar and a laser based radar (i.e. a lidar). Other sensing technologies will be considered at a later date.

This chapter reviews the current state of longitudinal collision avoidance and introduces the virtual bumper collision avoidance algorithm. An overview of the whole system will first be presented and then more detail will be given about the longitudinal controller.

1.1 Motivation

Since the advent of the automobile, engineers and scientist have sought to improve the safety of vehicles and of the driving task. Developments in the computer industry have made it possible to integrate embedded computing into products intended to lower the risk of injury due to a collision. Already common in today's automobiles are air bags, anti-lock brakes and computer designed impact absorbing chassis. These improvements have dramatically reduced the severity of crash related injury and fatalities.

Recently, research has been underway to use developing technology to actively prevent collisions as well as mitigate injury. This work falls under the area of Advanced Vehicle Control and Safety Systems (AVCSS), which aims to employ modern computing and sensing technology to assist the driver in the driving task. Some AVCSS applications are as simple as a warning for following too closely to the proceeding vehicle. Some advanced technologies can actually take control of the vehicle and make maneuvering decisions.

According to the 1997 General Estimates System (GES) crash database, there were 6.7 million police reported crashes. Of them, 37,280 resulted in fatalities [3]. The statistics on truck crashes are just as alarming. In 1997, 444,000 medium and heavy truck crashes resulted in 4,871 fatalities and 97,000 injuries nationally [3]. In addition to the significant loss of life, these accidents contribute to serious economic losses in terms of property damage and lost productivity. Our research aims to explore how technology can be used to prevent these considerable losses.

AVCSS can be broken down into two main areas of study: longitudinal control and lateral control. We will focus our review on the state of this technology on the longitudinal applications of AVCSS, especially the past several years, because this fits in well with the subject of this document. For an excellent review of the state of development of the whole AVCSS arena, see [4]. Schiller reviews longitudinal, lateral and combined longitudinal/lateral control in [2].

1.1.1 Longitudinal AVCSS

According to the 1997 GES crash database, there were approximately 1.921 million police reported rear end crashes in the U. S. This accounted for 28.4% of all police reported crashes that year [3]. The goal of longitudinal control systems is to reduce the number of rear end collisions.

Adaptive Cruise Control (ACC) was one of the first longitudinal control applications. ACC systems employ a range sensor (usually lidar or radar) that measures the distance and closing rate to the proceeding vehicle. Throttle and downshifting are used to bring the equipped vehicle to

the desired headway distance. In 1993, 26.5% of rear end collisions were caused by tailgating [5]. ACC systems can potentially reduce accidents by increasing the average headway to the preceding vehicle.

ACC has been the subject of considerable research for several years. One of the pioneers of this subject is Paul Fancher and his group at the University of Michigan [6]. They developed an ACC system on a Saab 9000 Turbo equipped with Leica's ODIN sensor. They conducted numerous field tests and analyzed the drivers' response to ACC. They have also compared headway times for manual driving, conventional cruise control and ACC.

Several car manufacturers have published their research results on ACC systems. BMW [7][8], Nissan [9], and Mitsubishi [10] have developed throttle control based ACC systems. This application is mature enough that within the year, Mitsubishi, Lexus and Mercedes-Benz plan to introduce ACC systems on several of their high-end models [11].

While the automakers have eagerly developed ACC and quickly brought it to the market, the development of ACC for heavy vehicles is less prevalent. One of the leaders in this area is Eaton VORAD. They have developed an ACC system for heavy-duty vehicles called SmartCruise™ [12]. This system was developed and installed on a Kenworth T600A tractor equipped with an electronically controlled Cummins diesel engine. The targets are sensed in front of the truck using a VORAD EVT-200 warning system. The velocity is controlled through throttle control and downshifting.

A new form of ACC called Enhanced Adaptive Cruise Control is the next step in complexity and functionality. Most ACC systems to date only employ throttle and gear shifting to decelerate the host vehicle, while Enhanced ACC offers "Low Speed Cruise" and "Stop & Go" capabilities [13]. Daimler Benz has developed a neural distance controller that employs braking [14]. The system was implemented on a MB 300 TE station wagon and can be used for ACC as well as for distance control in stop & go traffic. Only a few simulated driving maneuvers are necessary to

train the controller. Hitachi has developed an ACC system that controls the host vehicle speed with throttle, transmission and brake [15]. They developed a custom radar that provides range, range rate and azimuth angle to target. An automobile equipped with their system displayed the ability to decelerate 2.5 m/s^2 in response to a decelerating target vehicle. This level of deceleration is greater than can be achieved by throttle and gear shifting alone. The addition of braking makes this system able to respond to a greater variety of driving scenarios.

Collision warning systems provide the driver with an audio, visual or haptic signal when the host vehicle has reached a headway deemed unsafe by the collision warning algorithm. Eaton VORAD has developed a collision warning system for heavy vehicles [16]. It is sold as an off the shelf unit and contains a front radar and a side mounted radar for lane change warning (both sensors operate at 25 GHz). The warnings are given both audibly and visually.

Most collision warning systems reported in the literature come packaged with a collision avoidance system. One such system, developed by Lucas Advanced Engineering Center, was implemented on a Jaguar and contains ACC, collision warning and collision avoidance [17]. When the time to an obstacle becomes unsafe, a tactile warning is issued to the driver through a reactive force in the throttle pedal. If the rapid approach to an obstacle should become more critical, the vehicle automatically applies the brakes so as to mitigate or even avoid the impending collision.

The Daihatsu Motor Co. has developed a rear end collision warning/avoidance system that uses a visual display and automatic braking [18]. A head-up display shows the distance headway and warns the driver by blinking the master warning when there is a potential collision hazard. Brake actuation is then used if the driver does not respond.

Mazda Motor Co. has also developed a rear end collision avoidance system with automatic brake control [19]. This system uses scanning laser radar, path estimation algorithm with vehicle dynamics, collision prediction to the vehicle in front, and longitudinal automatic brake control.

The system demonstrated the ability to decelerate the test car to a stop in response to a stationary target. It also was able to apply braking to stop the host after the target vehicle applied emergency braking.

The final step up in complexity and functionality in the longitudinal AVCSS arena is full longitudinal control. Most applications of this sort involve a vehicle (or several vehicles) following the proceeding vehicle in a "master and slave" configuration. Usually, the vehicles are in communication using a radio frequency device and communicate their speed and/or acceleration.

One of the leaders in this area in the United States is the California Partners for Advanced Transit and Highway (PATH). Their system attempts to control the lateral and longitudinal motions of all vehicles in an automated multi-vehicle platoon system [20]. Sliding mode controllers were implemented on a four vehicle platoon with vehicle to vehicle communication. This four vehicle platoon demonstrated string stability at low to highway speed.

Another research effort in the U.S. was proposed by the Ohio State University. They are designing a truck convoy system that uses radar to detect a radar reflecting patch located on the proceeding vehicle [21]. The patch provides a strong return signal that can be easily discerned from noise so that the range, relative velocity, and angle to the proceeding truck can be measured. The radar system is augmented with vehicle to vehicle communications. The system has been tested in simulation with favorable results [22].

In Europe, the main longitudinal AVCSS effort was originally supported under the PROMETHEUS program. One particular European project sponsored by Daimler Benz called the Chauffeur project is focusing on the development of the 'tow-bar' concept. The idea involve the electronic linking of two (or more) commercial vehicles. The lead vehicle is driven by a human driver while the following vehicle(s) is electronically 'towed'. A vision-based sensor and vehicle to vehicle communication are used to provide the sliding mode headway controller

information on the proceeding vehicle [23][24]. Since they only have two properly equipped test vehicles, Daimler-Benz has successfully demonstrated the tow-bar concept with two vehicles. The addition of more vehicles to the platoon was tested by using data from previous test maneuvers as the sensor input to the test vehicles. Practical results of a seven truck platoon tested in this manner demonstrated string stability using the proposed control concept [25].

The University of Japan is researching longitudinal automated vehicle control of automobiles. They have tested two different control strategies on a two vehicle platoon. Sliding control and proportional integral differential (PID) control were implemented and compared. Sliding control showed better stability and more rapid convergence than PID control. Their vehicle models were found to be very accurate as the simulation and experimental results were impressively similar.

1.2 Overview of the Virtual Bumper

The virtual bumper is a collision avoidance strategy first introduced by Hennessey et al [26]. The virtual bumper combines longitudinal and lateral collision avoidance capabilities to control a vehicle in normal and emergency situations. This approach makes no assumptions about the dynamic capabilities of the surrounding vehicles in the driving environment when performing steering and throttle/braking actions. The dynamic limits of the host vehicle are considered but the control strategy does not guarantee that a collision will be avoided in all driving situations. In normal driving and some critical driving scenarios a collision can be avoided, but in the event of an unavoidable collision, the damage is mitigated.

The virtual bumper is a two dimension collision avoidance strategy for highway vehicles. The longitudinal subsystem can be thought of as a longitudinal collision avoidance system for heavy vehicles. We will demonstrate the longitudinal controller's ability to operate in an ACC mode, collision avoidance mode and in stop and go traffic. We did not come across any published literature of this type of system. The full range of operability of this part of the virtual bumper collision avoidance scheme for heavy vehicles makes it unique.

1.2.1 Personal Space and Region of Interest (ROI) Boundaries

A personal space is defined around the host vehicle. This personal space specifies the boundaries within which control action will be taken (see Figure 1.1). Vehicles inside this boundary impart a 'virtual' force upon the host vehicle, causing the vehicle to react to avoid collisions. We also define a road force that tends to keep the host laterally centered within its lane and away from the shoulder and road centerline.

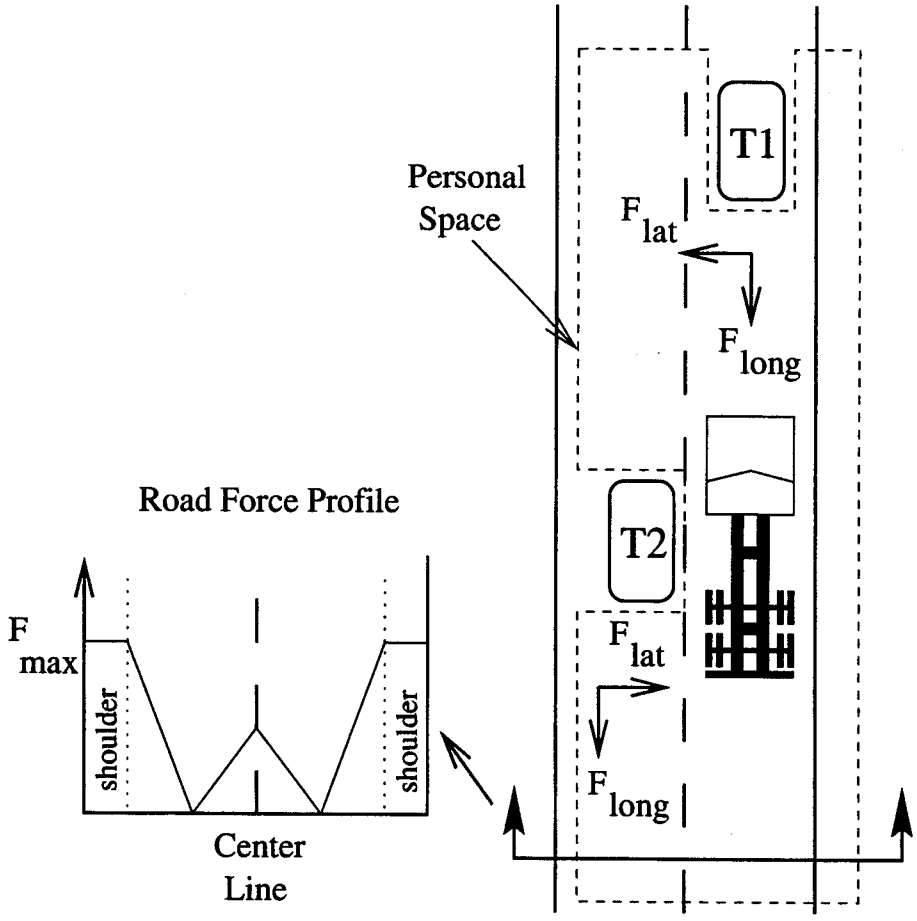

Figure 1.1: Overview of virtual bumper concept

The personal space is a subset of the region of interest (ROI, see Figure 1.2). The ROI boundary is determined by the maximum range of the host vehicle's sensing capability. A heuristic lane

change system uses the location and velocities of target vehicles in this area to determine the best lane for the host to travel. Vehicles in the ROI but not in the personal space do not impart a control virtual force on the host but do induce a lane changing force which is used for lane changing decisions only. Vehicles in the personal space (and ROI) cause both a control virtual force on the host and a lane changing force used for lane changing decisions.

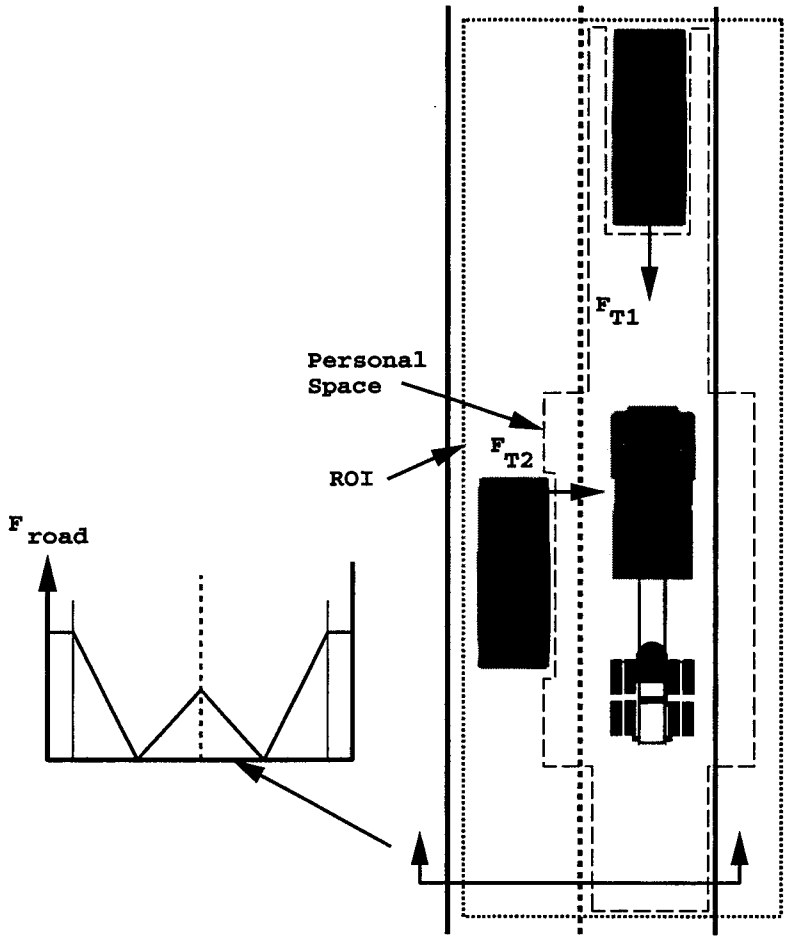

Figure 1.2: Personal space and region of interest boundaries

1.2.2 System Overview

The input into the virtual bumper is the current state of the host vehicle and the location and trajectory of the vehicles surrounding the host. The output is a desired traveling speed and a

lateral lane location/offset (Figure 1.3). The algorithm consists of three main modules; the longitudinal controller, lateral controller and lane change system. The longitudinal controller determines the desired host speed by applying forces in the longitudinal direction. The lateral controller produces distance offsets in the lateral direction as a reaction to vehicles in the personal space. The lane change system determines the best lane of travel based on all the target vehicle locations and relative velocities in the ROI. When the current lane of travel is chosen, the lateral position is determined by the lateral controller. When the current lane of travel is no longer the safest lane of travel, the lane change system produces lateral displacement commands that generate a lane change trajectory.

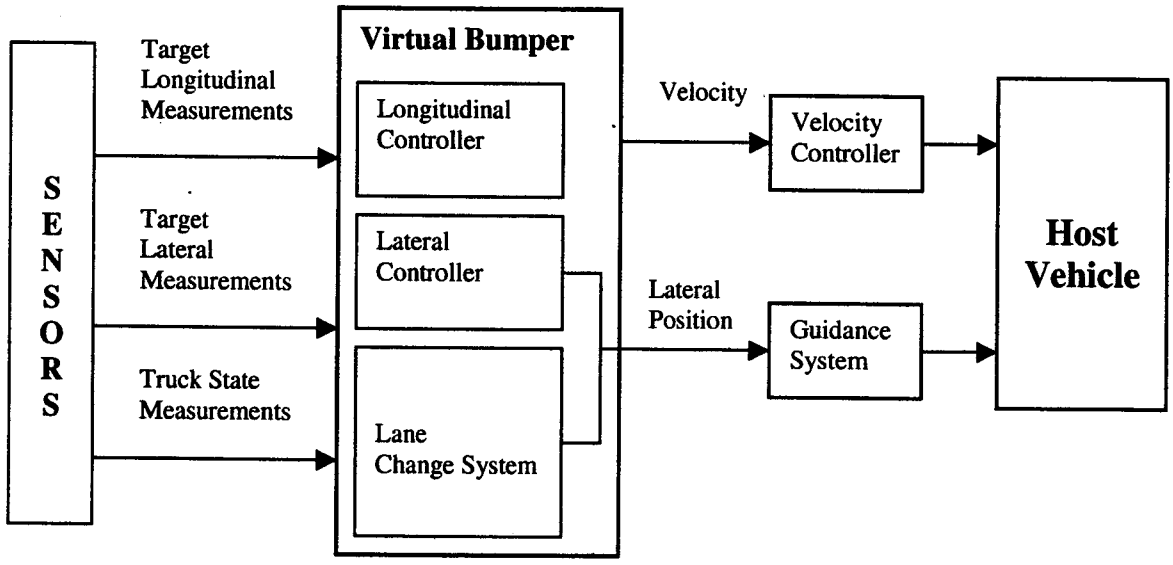

Figure 1.3: General system block diagram of the virtual bumper

The virtual bumper can be thought of as a higher level control system. The velocity controller and guidance system are subsystems that attempt to move the truck's state to one that the virtual bumper has commanded. Therefore, the virtual bumper does not directly control the movement of the vehicle. Instead, it sends velocity and lateral position commands to the velocity controller and navigator subsystems, respectively. An assumption is made that the lower level controllers are able to follow the commands given to it from the virtual bumper. Subsystem time constants

and host vehicle dynamic limits are considered to ensure that the virtual bumper output commands are reasonable.

1.2.3 Sensing the Road Scene

It is clear that the virtual bumper must know the location and trajectory (relative velocities) of the vehicles in the environment surrounding the host. This information may be gathered using various sensing strategies. Schiller [2] developed a range sensor strategy in simulation which effectively gathered the desired target information in the host's surroundings. The strategy employed a radar sensor array that surrounded the host vehicle. Each simulated radar sensor was adjustable in location and orientation around the perimeter of the host vehicle. An optimal configuration involving 14 radar sensors was developed and used as the sensing input to the virtual bumper algorithm.

While Schiller was successful in using an array of radar sensors to sense the surrounding host environment in simulation, a real world implementation poses several difficulties, which may make this strategy difficult to realize. The sensor array strategy proved to be sensitive to radar mounting orientation. Practical mounting error (tolerance) may be enough to cause improper lane determination since the strategy used the radar's orientation and beam geometry to determine in which lane the target resides. Another practical issue with this strategy that was not evaluated in the simulation is cross talk. Several radar units located next to each other may receive reflections from neighboring sensors and produce invalid target signals due to crosstalk. The final practical limitation of this approach is cost. Given the current cost of radar technology, a radar sensor array consisting of fourteen radar units may be cost prohibitive in a practical deployment. But costs are dropping and may continue to do so making sensor arrays a reasonable approach, especially if they have a wider field of view and the capability to locate the azimuth of an obstacle in addition to range and closing rate.

For future testing of the virtual bumper, we feel a combination of radar and Differential Global Positioning System (DGPS) will be a practical solution. Vehicles can transmit their location and

trajectory to each other on a moving Local Area Network (LAN). The host vehicle will then know the location of all vehicles present within the maximum range of the transmission device. Radar can be used on the front of the host for fast detection of vehicles in the immediate path of the host and for sensing vehicles or obstacles not capable of transmitting their own position. Redundancy in this region is desirable.

Once the algorithm is fully realized and tested using vehicle to vehicle communication, experiments may be performed to simulate different radar configurations based on the latest specifications of modern commercially available radar products. We will also investigate the possibility of custom radar units (some of which are quite inexpensive) based on the results of the vehicle to vehicle communication experiments. At that time, mounting issues and cross talk will be addressed. Other sensing modalities can also be investigated in the future.

1.2.4 Test Bed

The virtual bumper demonstrated acceptable collision avoidance capabilities in the simulations performed by Schiller [2]. We plan to implement the system on real hardware in three stages. This thesis documents the implementation and testing of the first stage, the longitudinal controller. The lateral controller and the lane change system will be implemented later.

One of the reasons that we chose to implement the longitudinal controller in the first phase is due to the existing sensing technology aboard the Navistar tractor test bed. Our sensing platform currently consists of a rate gyro, magnetometer, accelerometers, DGPS, a radar and a laser range finder. We do not currently have a means of sensing target vehicles in all directions around the tractor. We do plan on equipping the tractor and at least two cars with wireless communication devices to form a LAN in which the vehicles can communicate their state information. We used an Eaton VORAD EVT-200 radar and an Optech Sentinel 100 range finder as the target sensors for the implementation of the longitudinal controller, which only needs sensor coverage in the front of the host vehicle.

As stated previously, the goal here was to test the control algorithms and not to determine the optimal hardware configuration. Results obtained in this phase will be directly applicable when the final sensing strategy is selected. The software was designed so that changing any of the sensing hardware will require re-writing only one software module. The interface to the virtual bumper is not sensor specific.

1.3 The Range vs. Range Rate Phase Diagram

Before we introduce the control theory of the longitudinal controller, an understanding of the range vs. range rate phase diagram is essential. This diagram, introduced by Fancher et al [6] [27], is a plot of the range (R) versus the closing rate or range rate ($\dot{R}$) to the proceeding vehicle. This diagram graphically describes the relationship between the host and target's velocities and spacing. Thus, this diagram can be used to design and evaluate headway controllers. Additionally, regions in this diagram can be defined to describe real driving situations (i.e. braking required, collisions imminent, etc.). In this section, we will briefly describe the R, $\dot{R}$ phase diagram so that the reader understands the design of the headway controller. We will also use this diagram to evaluate experimental data and determine how well the control meets the design specifications. More details for the interested reader can be found in [28].

1.3.1 Basic Concepts of Headway Control

The objective of a headway controller is to control the distance at which the proceeding vehicle is followed. The headway distance can be defined in various ways: constant headway distance, constant time headway or adaptive time headway. We chose to define a constant headway time and then derive the desired headway distance based on velocity. The headway time can be thought of as a velocity based measure of how close the host follows the target. Lower headway time means smaller headway distances and vice versa.

The desired headway is determined by the headway time and the velocity of the target vehicle (Equation 1.1). It is not based on the host vehicle velocity because in order to arrive at a steady state headway distance, the host must change its velocity to match that of the target vehicle.

$$R_H = T_H V_{target} \tag{1.1}$$

Headway controllers also monitor the speed differential between the host and target and attempt to make this value zero once the headway is at the desired state. Radar is used to measure the speed differential using its range rate (or range dot) measurement (Equation 1.2).

$$\dot{R} = V_{tar} - V_{host}, \tag{1.2}$$

where

V_{tar} = velocity of the target vehicle, and

V_{host} = velocity of the host vehicle.

From the previous equations it is apparent that the relationship between range and range rate plays an important role in headway control. We used this fact to design the headway controller of the virtual bumper.

1.3.2 R vs. $\dot{R}$ Trajectories

The range vs. range rate phase plot is formed by placing range measurements on the y-axis and range rate on the x-axis. Figure 1.4 shows a scenario in which the host vehicle is traveling toward a slower moving target. A negative range rate means that the host vehicle is approaching the target vehicle. The range can only be positive (it is zero when a collision has occurred) so all headway control occurs in the upper half of the plot. The figure shows a constant deceleration curve as the host decelerates until the range rate is zero and the desired headway is reached. Constant deceleration trajectories in range vs. range rate phase plot are parabolic (the equation in Figure 1.4 will be derived in §1.4.3)

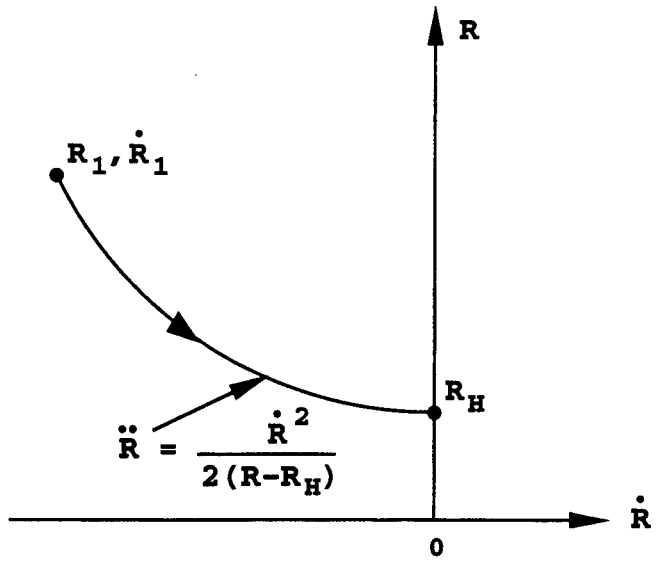

Figure 1.4: Constant deceleration trajectory on R vs. $\dot{R}$ phase plot

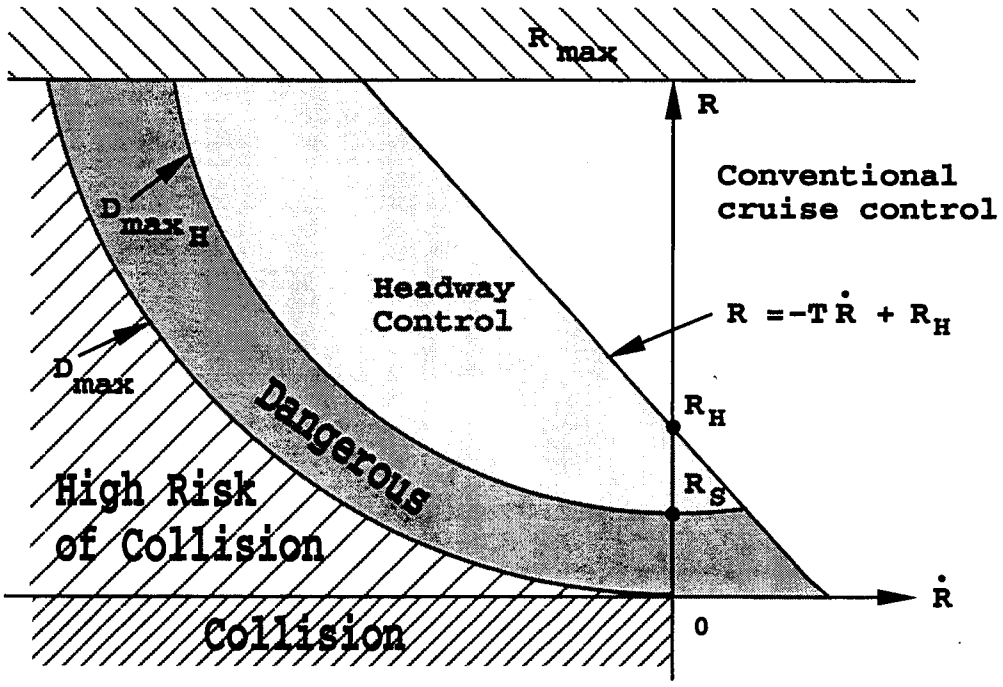

Figure 1.5: Zones in the range vs. range rate phase plot

We can define zones in the phase plot as shown in Figure 1.5. The R_{max} region is the region beyond the maximum sensing capabilities of the host based sensors. Target vehicles above this zone are not detected. We define a line that determines when the headway controller is active as follows:

$$R = -T\dot{R} + R_H. \tag{1.3}$$

The time constant (T) is determined by the gains of the impedance controller, which we will discuss shortly. To the right of this line is conventional cruise control. In this region, targets are not in the personal space of the host so the velocity controller aims to achieve the user set cruising speed. To the left of this line the headway controller is active. The constant deceleration curve $Dmax_h$ is the maximum deceleration capability of the headway controller, which is determined by the maximum deceleration capability of the velocity controller. The constant deceleration curve ($D_{max)}$ is the maximum deceleration capability of the host vehicle. The region to the left of this line represents a situation where the host vehicle does not have the deceleration capability to avoid a collision with the target. A collision will ensue unless the target increases its velocity or either vehicle engages in a lateral evasive maneuver. The range cannot be a negative value, but the collision area indicates that if the range goes to zero and the range rate is negative, a collision has occurred.

1.4 Longitudinal Controller

The longitudinal controller controls the headway distance to the preceding vehicle as well as maintains the desired velocity while operating under conventional cruise control. The headway controller is the component of the longitudinal controller that maintains the desired headway to the preceding vehicle. The second part of the longitudinal controller is the cruise control impedance, which acts to maintain the desired traveling speed when the target exits the personal space (conventional cruise control).

1.4.1 Longitudinal Controller Block Diagram

The longitudinal controller block diagram is shown in Figure 1.6. The range and range rate of the target vehicle are sensed by the range sensors and are inputs, along with the host velocity, to the longitudinal impedance block. The output of this block is the virtual longitudinal force generated by the "compression" of the target's virtual bumper impedance that acts on the admittance of the vehicle. This admittance is related to the vehicle's actual dynamics and represents the possible velocity change that can be generated by the virtual forces. A velocity change is produced and added to the reference velocity. The command velocity is sent to the velocity controller which produces throttle/brake commands to "force" the host vehicle to the desired velocity. SW1 and SW2 are software switches that are never closed at the same time. The switch SW1 is turned on when a target is with the longitudinal personal space. The switch SW2 is turned on when there is no target in the personal space.

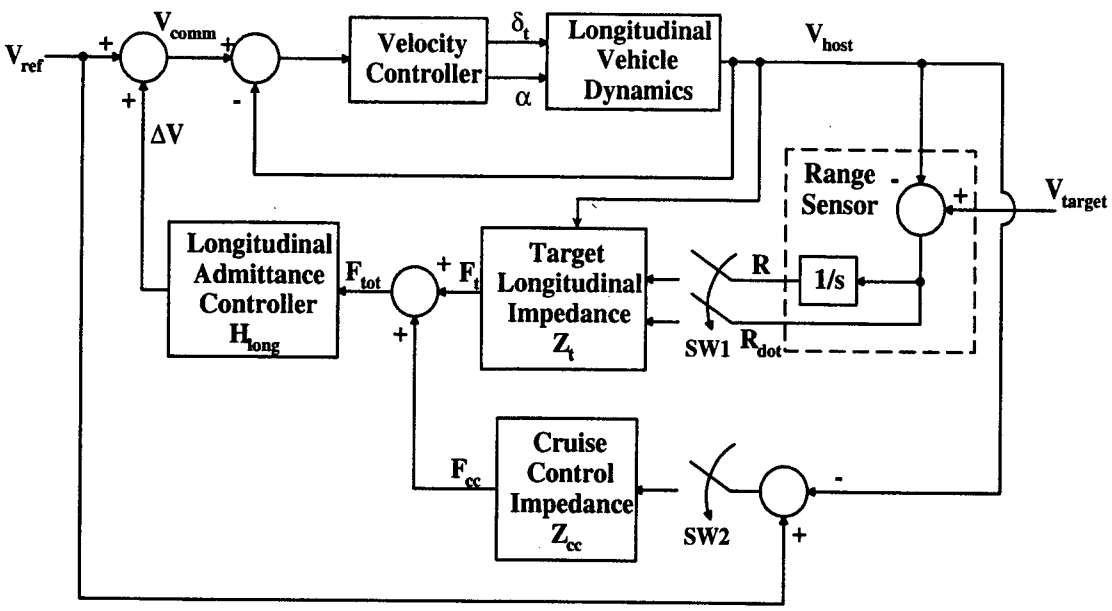

V_{ref} = Operator selected nominal velocity

ΔV = Offset in desired velocity calculated by impedance controller

V_{comm} = Command velocity

δ_t = Commanded throttle position

α = Commanded brake position

V_{host} = Host vehicle velocity

V_{target} = Target vehicle velocity

s = Laplacian operator

R = Range to target measured by radar sensor

R_{dot} = Range rate to target measured by radar sensor

SW1 = Software switch that is enabled when target is within ROI

SW2 = Software switch that is enabled when target is not within personal space

Z_t = Virtual impedance associated with space between target and host vehicles

Z_{cc} = Virtual impedance associated with difference between reference and host velocity

F_t = Virtual force calculated for target which is applied to host vehicle

F_{cc} = Virtual cruise control force

H_{long} = Longitudinal admittance controller

Figure 1.6: Longitudinal control block diagram

Linear and nonlinear forces are generated by the longitudinal bumper impedance. Linear forces are defined by a linear response on the range vs. range rate diagram. This type of force is used when minimal braking is required to reach the desired headway. When greater levels of braking are needed, a non-linear force is used that has a parabolic trajectory on the range vs. range rate phase plot. This is a constant deceleration force.

We will first introduce the linear force and the impedance control law which governs its behavior. We will then discuss the non-linear force and explain when each of these forces is used.

1.4.2 Definition of Linear Longitudinal Force and Personal Space Boundary

The longitudinal personal space extends from the front bumper of the host vehicle and between the lanes of travel (Figure 1.7). The personal space free length is a function of the speed of the target vehicle as well as the headway time. A virtual spring (with a spring coefficient of k) and a virtual damper (with a damping coefficient of b) make up the impedance controller. Deflections in the spring free length cause a virtual spring force. Range rate values (relative velocity) determine the virtual force of the damper as shown in Equation 1.4.

$$F_{long\_lin} = k(R - R_H) + b\dot{R} \qquad (1.4)$$

The personal space free length (R_{PS_lin}) is calculated by setting the longitudinal force in equation 1.4 to zero and plugging it in equation 1.1 to arrive at the following:

$$R_{PS\_lin} = -(T_H + b/k)\dot{R} + R_H . \qquad (1.5)$$

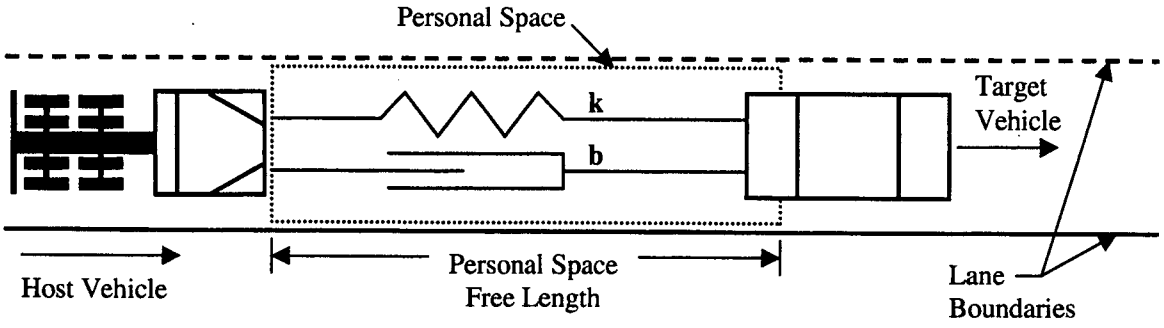

Figure 1.7: Longitudinal controller personal space boundary

The personal space free length forms the switching line in the range vs. range rate phase diagram as shown in Figure 1.8. This line represents the range at which to activate the headway controller and has a slope that is determined by the desired headway time (T_H, set by the user) and the spring and damper coefficients (b and k, respectively).

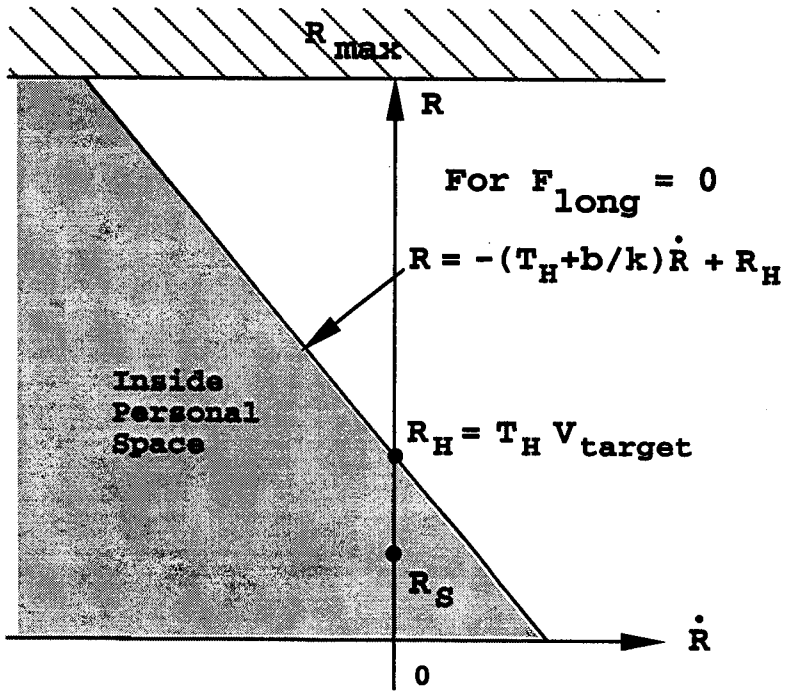

Figure 1.8: Switching line that determines when headway control is active

The characteristic equation for this system with the linear force defined above is described by a second order expression in the Laplace domain given by Equation 1.6:

$$Ms^2 + bs + k \tag{1.6}$$

This equation determines the dynamic response characteristics of the host vehicle's headway as it approaches a target vehicle. Choosing the proper values of b and k will produce a desirable response.

We should first define what is a desirable response. Clearly, the headway response should be non-oscillatory with little or no overshoot so that the deceleration is comfortable. Furthermore, the linear command response should be slow enough so that minimal braking is required. Large acceleration levels are the responsibility of non-linear forces, which will be defined shortly.

To meet the design criteria, the gains should be chosen so that the system is over-damped. The damping ratio (ζ) and the natural frequency (ω_n) are given by the following characteristic expression:

$$s^2 + 2\zeta\omega_n s + \omega_n^2 \tag{1.7}$$

Equating coefficients of Equations 1.6 and 1.7 yield,

$$\frac{b}{m} = 2\zeta\omega_n \tag{1.8}$$

$$\frac{k}{m} = \omega_n^2 \tag{1.9}$$

Although the characteristic equation is second order, it can be made to approximate a first order response by setting the damping (b) and spring stiffness (k) coefficients such that the system is over-damped. An over-damped system has two real negative poles (solutions to the

characteristic equation). The pole closer to the origin is the dominant pole (p_1) and has the greatest effect on the dynamics of the controller.

Factoring yields

$$(s+p_1)(s+p_2) \tag{1.10}$$

where

$$p_1 = -\omega_n\left(\zeta - \sqrt{\zeta^2 - 1}\right), \text{ and} \tag{1.11}$$

$$p_2 = -\omega_n\left(\zeta + \sqrt{\zeta^2 - 1}\right). \tag{1.12}$$

After the dominant pole is located, the second pole can be obtained by selecting the damping ratio (ζ) greater than or equal to one (over-damped).

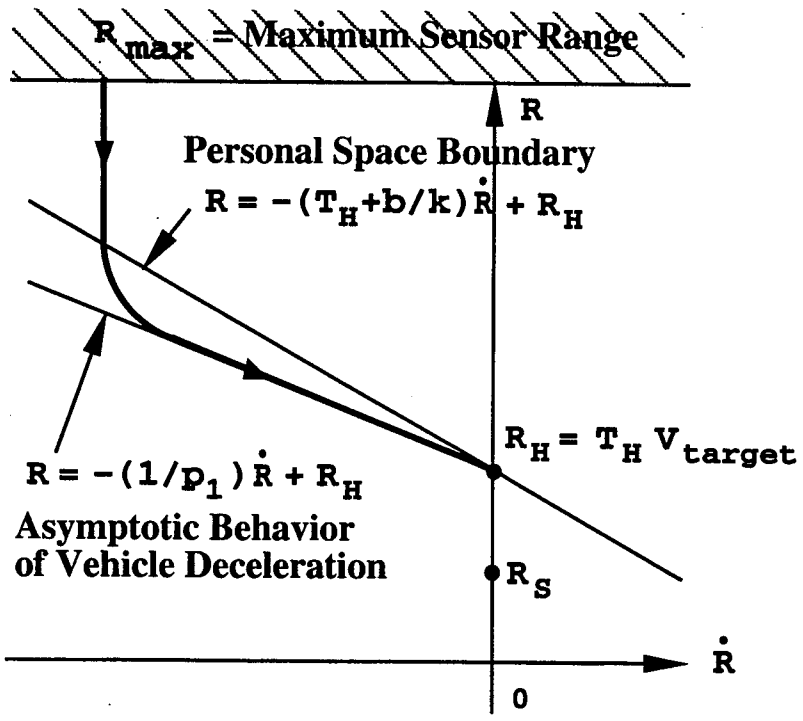

Figure 1.9: Dominant pole response of impedance controller

Figure 1.9 shows the expected headway response determined by the dominant pole. The range decreases as the host approaches the target and the range becomes less than the personal space boundary. The controller then issues force commands that smoothly slow the host vehicle to match the velocity of the target at the desired headway, R_H. R_S is the minimum safe headway distance and is used in non-linear force calculations (section 1.4.3).

1.4.3 Definition of Nonlinear Longitudinal Force and Personal Space Boundary

When greater deceleration levels are required to achieve the desired headway, non-linear forces are used. We call this force non-linear in reference to the effect this force has on the headway response in the range vs. range rate phase plot. While the linear force was designed for a 'straight line' behavior in this domain, the non-linear force was designed to have a parabolic curve. This is a curve of constant deceleration on the range vs. range rate phase plot.

Consider a scenario in which the host is traveling 55 mph and the target is moving 15 mph. The linear force equation would reduce the host vehicle velocity in roughly a first order decay type of response. This would cause large acceleration in the beginning of the maneuver, after which only small accelerations are needed to smoothly track the target vehicle's velocity. Clearly, this would not be comfortable for the driver and does not simulate a human driver's response to such a scenario. A more suitable response would be a constant deceleration down to the desired headway with little or no overshoot. This type of response can be achieved with non-linear forces that produce constant deceleration.

The equation for constant deceleration levels can be derived from the time domain equations for R and $\dot{R}$.

$$R(t) = R_1 + \dot{R}_1 t + \frac{1}{2}\ddot{R}t^2 \tag{1.13}$$

$$\dot{R}(t) = \dot{R}_1 + \ddot{R}t \tag{1.14}$$

Since at the goal point, $\dot{R} = 0$, equation 1.14 is set equal to zero and solved for t so that,

$$t = -\frac{\dot{R}_1}{\ddot{R}}. \tag{1.15}$$

Plug this value for time into equation 1.13 and solve for the deceleration level,

$$D_{host} = \ddot{R} = \frac{\dot{R}_1^2}{2(R - R_H)}. \tag{1.16}$$

This equation clearly shows that constant levels of deceleration are represented by parabolas in the range vs. range rate phase plot.

One thing to notice about Equation 1.16 is that as the target vehicle decelerates and the range decreases, the range approaches the desired headway (R_H). This causes Equation 1.16 to approach infinity. To prevent this from occurring, R_H is replaced with a scaled headway, $R_{scaling}$, which is a range between R_H and a minimum safe headway R_S. Then the value of $R_{scaling}$ is determined by the deceleration level calculated for the previous time step as,

$$R_{scaling} = R_S + (R_H - R_S)\left(\frac{(D_{host}(n-1) - D_{PS\_nlin})}{(D_{max} - D_{PS\_nlin})}\right) \tag{1.17}$$

where,

R_S = Minimum safe headway

D_{max} = Maximum deceleration capability of the host

D_{PS_nlin} = Deceleration level which defines the personal space for the nonlinear force

$D_{host}(n-1)$ = Calculated deceleration level at last time step

The minimum safe headway is calculated in the same manner as the desired headway based on a headway time T_s, such that

$$R_S = T_S V_{target}, \tag{1.18}$$

where

T_S = Minimum safe headway time.

The non-linear force personal space boundary is established by setting a deceleration level under which non-linear forces are applied. This deceleration level should be small enough so that minimal braking levels are applied in the linear force region.

As previously stated, constant deceleration produces a parabolic curve in the range vs. range rate phase plot. Non-linear forces will be applied below and to the left of this curve, which is determined by plugging D_{PS_nlin} into Equation 1.16. Solving for the range (R_{PS_nlin}) yields,

$$R_{PS\_nlin} = R_S + \frac{\dot{R}^2}{2(D_{PS\_nlin})} \tag{1.19}$$

The region where each force type is applied is shown in Figure 1.10. To the right of the linear personal space switching line no virtual forces are applied to the host. To the left of the linear personal space switching line and to the right of the non-linear personal space curve, linear forces are applied. To the left of the non-linear personal space curve, non-linear forces are applied. The maximum deceleration curve (D_{max}) shows the maximum deceleration that can be generated by the host vehicle. The non-linear force equation is finally determined by multiplying the deceleration equation by the mass of the vehicle (M),

$$F_{long\_nl} = \max\left\{\begin{array}{c} -\dfrac{M\dot{R}^2}{2(R - R_{scaling})} \\ \\ -MD_{max} \end{array}\right\} \tag{1.20}$$

The non-linear longitudinal forces must not exceed the maximum deceleration limits of the velocity controller. Thus, the greater of the two forces is chosen.

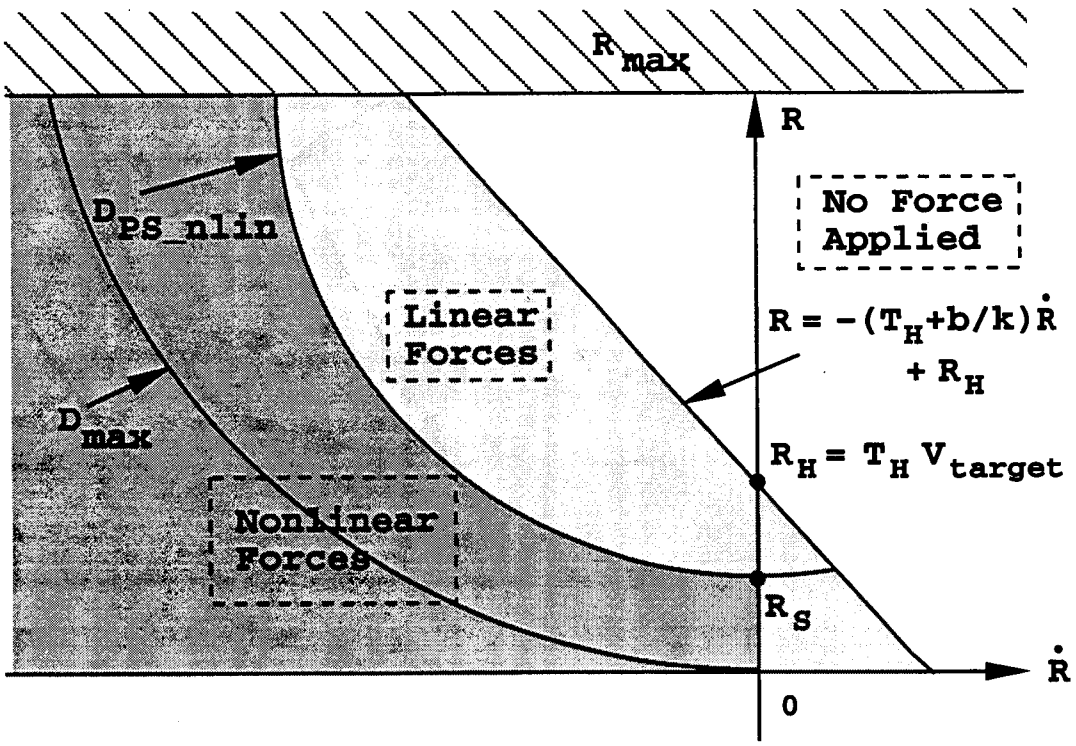

Figure 1.10: Force regions in the R vs. $\dot{R}$ range rate phase plot

To better understand the switching between linear and non-linear forces, we will go through an example. In this example, the host is traveling at a significantly higher speed than the target vehicle such that when the range sensor detects the target, it is at a point in the non-linear region of the range vs. range rate plot (Figure 1.11). Non-linear forces are applied to the host, which cause a constant deceleration until the range and range rate reach the switching line between the linear and non-linear force region (determined by D_{PS_nlin}). The forces switch to linear which cause the trajectory of the target state $(R, \dot{R})$ to change. The host vehicle slows until it matches the target vehicle's velocity. This happens at a range that is less than the desired headway. Thus, the host slows to a velocity less than the target until the headway increases to R_H.

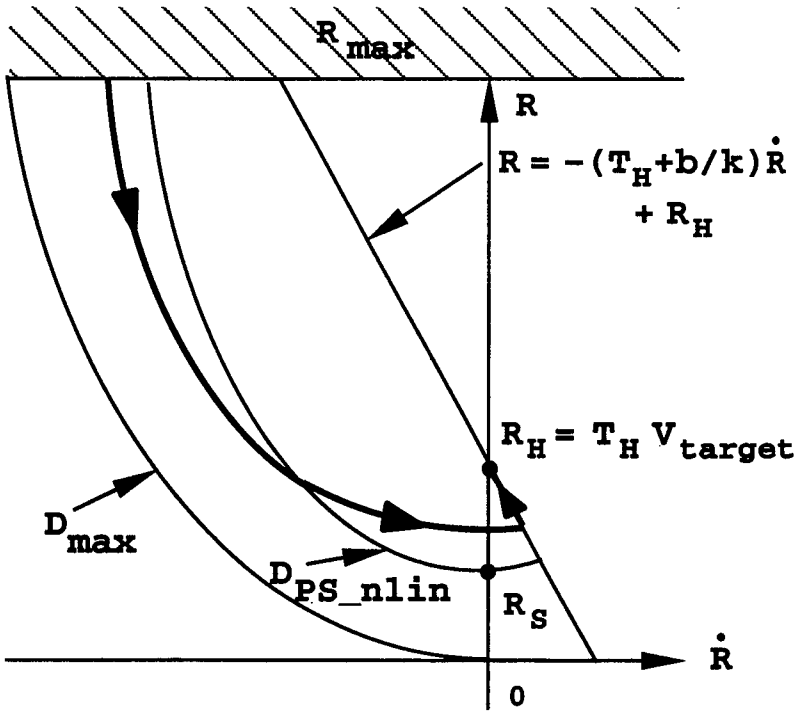

Figure 1.11: Example of switch from non-linear to linear forces

1.4.4 Host Vehicle Admittance

The virtual force acting on the host vehicle is only one of several that determine the host's velocity. The equation of motion on a vehicle consists of several forces acting on the mass of the vehicle as shown in Equation 1.21 [29]:

$$ma = F_{engine} - F_{brake} - F_{roll} - F_{aero} - F_{grade}, \tag{1.21}$$

where

F_{engine} = Engine torque transferred to a force by the transmission and tires,

F_{brake} = Braking torque transferred to a force by the tires,

F_{roll} = Rolling resistance of the tires,

F_{aero} = Aerodynamic resistance to motion,

F_{grade} = Force induced by grade of road.

The relationship between the force applied to this system by the throttle/brake and the resulting velocity is approximately first order. The transfer function that inputs force and outputs velocity is given by

$$H_{long}(s) = \frac{V(s)}{F(s)} = \frac{1}{ms + b}, \tag{1.22}$$

where

m = the mass of the vehicle,

b = the truck damping coefficient (not to be confused with the damping coefficient of the virtual damper) which is a function of the velocity, rolling resistance, aerodynamic resistance and road grade changes.

While this transfer function is accurate, it is not easy to derive because extensive experiments must be performed to obtain the damping force function. Thus, we will neglect this term (b) so that the host vehicle admittance is

$$H_{long}(s) = \frac{1}{ms} \tag{1.23}$$

This admittance produces velocity offsets that contribute to the desired velocity to produce the needed command velocity. The command velocity is then sent to the velocity controller, which adjusts the throttle and brakes. Since the longitudinal forces are limited by deceleration limits and impedance controller gains, the velocity offsets should be reasonable.

1.4.5 The Cruise Control Force

The final component of the longitudinal controller is the cruise control impedance force. This force is required because the host may track a slower moving target vehicle that switches lanes. The forces are zero when no target is located within the longitudinal personal space. However, the velocity offset from the admittance part of the controller would keep the velocity command the same as the last velocity command given before the target switched lanes. A cruise control

impedance was added to ensure that the host achieves the desired traveling velocity when no target is located in the longitudinal space.

It should be noted that the role of cruise control impedance could have been carried out by resetting the admittance control block output to zero once the target leaves the personal space boundary. This would effectively produce a zero velocity change command from the virtual bumper so that the user set desired velocity is inputted to the velocity controller. We chose to use a cruise control impedance/force because it is more consistent with the overall approach and it is more robust to range sensor dropouts.

The cruise control force is generated when no target is in the personal space. This force acts on the velocity difference between the actual and user selected cruising speed. The cruise control force executes the standard cruise control functionality of the longitudinal controller.

1.5 Implementation Issues

As explained previously, the longitudinal controller will be implemented first with the existing sensor technology mounted on the SAFETRUCK. (The SAFETRUCK is an International Navistar 9400 equipped with computers and DGPS, IMU and radar sensors.) The target sensing hardware consists of both a microwave based radar and a laser based range finder. These sensors were installed in the front center location of the host so that their coverage area coincides with the longitudinal personal space.

In deriving the impedance controller we made an assumption in the previous section. We assumed that the velocity controller was able to track velocity commands with little error or latency. This, of course, is not the case. The velocity controller has its own time constant and we must compensate for the tracking error by modifying the linear and non-linear force equations.

Recall the linear force equation (repeated here in Equation 1.24) that was designed in order to make the headway approach the desired value with a smooth first order response.

$$F_{long\_lin} = k(R - R_H) + b\dot{R} \tag{1.24}$$

The tracking error of the velocity controller would cause the headway to be oscillatory and causes over-shoot in the system response. This is due to the delay in the velocity controller achieving the commanded velocity. The virtual forces increase because the velocity of the host has not immediately changed as expected. The continued application of the virtual force (until the host vehicle's velocity changes) causes the host vehicle's velocity to decrease to a level below what was needed. The virtual force then become positive and the velocity oscillations continue.

To eliminate this behavior, we add predictive terms to the force equations that use the present system state to predict the future system state. We replace the terms R, $\dot{R}$ and R_H in the force equation with the following predictive terms.

$$R_{pred} = R + \dot{R}T - \frac{1}{2}\dot{V}_{host}T^2, \tag{1.25}$$

$$\dot{R}_{pred} = \dot{R} - \dot{V}_{host}T, \text{ and} \tag{1.26}$$

$$R_{H\_pred} = T_H\left(V_{host} + \dot{V}_{host}T\right), \tag{1.27}$$

where

T = Time constant matched to the velocity controller.Substituting the predictive equations into Equation 1.24 leads to the modified linear force equation shown in Equation 1.28.

$$F_{long\_lin}^{'} = k(R - R_H) + (kT + b)\dot{R} - \left(k\left(\frac{T^2}{2} + T_H T\right) + bT\right)\dot{V}_{host} \tag{1.28}$$

The predictive time constant (T) must be tuned to eliminate the oscillation caused by the tracking error of the velocity controller. Notice that the modified linear force equation now has a feedback term that includes the host acceleration. This term dampens the control response as a means to limit rapid acceleration changes. This may slow the response of the controller. However, this linear force equation is only applied to situations where little or no braking is required. The advantage of a smooth non-oscillatory response outweighs the slower response, especially since the non-linear forces act as a backup when a faster response and higher forces are needed.

Since the linear force equation has changed, the linear personal space boundary has also changed. We must recalculate this by setting Equation 1.28 equal to zero and solving for the range. Note that the initial acceleration of the host vehicle is zero when traveling at the desired steady state velocity.

The equation

$$R_{PS\_lin} = T_H V_{target} - \left(T + \frac{b}{k}\right)\dot{R}, \qquad (1.29)$$

reveals that the personal space boundary increases by $T\dot{R}$ due to the predictive time constant, implying that the impedance controller will activate at longer range when $\dot{R}$ is negative.

The nonlinear force equation must also be modified to compensate for the tracking error of the velocity controller. We repeat the equation for non-linear forces below.

$$F_{long\_nl} = \frac{M\dot{R}^2}{2(R - R_{scaling})} \tag{1.30}$$

Just like in the linear force equation, we compensate for the tracking error in the velocity controller by replacing the range, R, with the predicted range as shown in Equation 1.31:

$$R'_{pred} = R + \dot{R}T. \tag{1.31}$$

Notice that this is slightly different than the predictive term in the linear force equation in that the host acceleration ($\dot{V}_{host}$) is not included. This was done because the acceleration feedback reduces the forces in order to reduce oscillations and over-shoot. Nonlinear forces are used when large amounts of braking are needed so this term would reduce forces, which is not desirable. Moreover, the nonlinear forces were designed so that they do not have oscillations or overshoot due to the parabolic constant deceleration profile in the range vs. range rate phase plot. Thus, the acceleration feedback term is not needed. Substituting Equation 1.31 into Equation 1.30 results in the following non-linear force equation:

$$F'_{long\_nl} = \frac{M\dot{R}^2}{2(R + \dot{R}T - R_{scaling})}. \tag{1.32}$$

No modification to the non-linear personal space (R_{PS_nl}) is needed because this boundary is not based on solving for the range that solves Equation 1.33 when F_{long_nl} is equal to zero. Rather, this boundary is determined by a constant deceleration limit, below which non-linear forces are applied. Therefore, the non-linear personal space is unchanged from equation 1.19.

A slight modification in the definition of desired headway (R_H) was also made to facilitate a real world implementation. The virtual bumper is a collision avoidance scheme that should be robust for all driving situations. Consider a scenario where the host vehicle approaches a traffic jam. If the traffic is at a stand still, the desired headway would be zero ($V_{target} = 0$) and the controller

would attempt to reach this headway. Clearly, it would not be safe or comfortable for the drivers of the host and target vehicle to be at such a close proximity. Therefore, we have modified the desired headway calculation to add a minimum safe headway as follows.

$$R_H = T_H V_{target} + R_{Ho} \tag{1.34}$$

where

T_H = Desired headway time, and

R_{Ho} = Desired headway at zero velocity.

Similarly, the minimum safe headway (R_S) must be modified so that it does not approach zero, i.e.,

$$R_S = T_S V_{target} + R_{So}, \tag{1.35}$$

where

T_S = Safe headway time, and

R_{So} = Safe headway at zero velocity.

We have defined all the major components of the virtual bumper and have in detail defined the control law of the longitudinal controller. In the next chapter, we set the time constants of the longitudinal controller and first implement it in simulation to determine the effect of the controller gains on the host vehicle.

Chapter 2

Simulation Results

2. Simulation Results

In this chapter we select the parameters for the virtual bumper force equations based on a computer simulation of the longitudinal controller and the modeled dynamics of the International Navistar 9400 tractor. The simulation is run with various headway controller gains to give the reader a feel for their effect on the host-target vehicle interaction as well as to determine a starting point for the actual implementation. Sensor noise and latency are also simulated to determine how they affect the algorithm. By the end of this chapter the reader should have a good feel for some of the issues involved in the actual implementation based on what was discovered in the simulation runs.

2.1 Parameter Selection

The previous chapter provided an overview of the virtual bumper collision avoidance scheme. Specifics were provided on the control laws that govern the longitudinal portion of the virtual bumper. Several parameters were presented which affect the behavior of the algorithm. In this section we will determine the values of certain parameters and explain the reasons for the selections.

One of the most important parameters in the system is the headway time, T_H. This variable ultimately decides the desired headway based on the velocity of the target vehicle. Fancher has demonstrated that desired headway time differs based on many factors. Age and personality type seem to play the biggest role in a person's choice of a comfortable headway [30]. For this reason, we chose to let the headway time be selected in real time through a control panel. For the implementation experiments, we chose this value to be between 1.0 and 2.0 seconds, but this was arbitrary and can be changed by the user.

Since the virtual bumper does not guarantee an accident would not occur, the term "safe" headway is a little misleading. It was named as such to make sure that the non-linear forces do not decrease the headway below a minimum "safe" distance. The parameter used to set this distance is the safe headway time, T_S. We chose to set this at half the headway time, or, $T_S = T_H/2 = 0.5$ seconds.

2.1.1 Velocity Controller Deceleration Limits

To determine the maximum deceleration limit of the headway controller ($Dmax_h$) and the non-linear personal space deceleration limit (D_{PS_nl}) we must evaluate the velocity controller. The velocity controller was designed and programmed by Lee Alexander [31]. The longitudinal controller is reliant upon the velocity controller to follow its commands. To ensure that the virtual bumper produces velocity commands that can be followed by the velocity controller, we set limits on the maximum allowable forces to match the dynamic capabilities of the velocity controller. To that end, we performed deceleration experiments to test the deceleration capabilities of the velocity controller and Navistar tractor.

The first experiment was designed to determine the deceleration of the Navistar tractor without a trailer, due to rolling and aerodynamic resistance forces. This allowed us to determine at what deceleration level braking is applied. The Navistar tractor was given a constant command velocity of 15 m/s. When the truck reached this velocity, the command was stepped down to zero and the truck was allowed to coast. The braking system was inactive for this experiment (see Figure 2.1). Notice that the truck slowed down to around 3.5 m/s and then maintained this velocity. The velocity data appears fairly linear during the 'coasting' deceleration.

The linear section of the deceleration plot in Figure 2.1 is shown in Figure 2.2. A linear fit of the data is also shown and reveals that the slope of the fitted line (acceleration) is –0.37 m/s^2 (0.038 g's). The data has a very linear trend, confirmed by the R^2 value of 0.9955.

The 'coasting' deceleration is rather small, so in order to ensure that light braking would be included in the linear force region, we chose to set the non-linear personal space boundary, D_{ps_nl}, at 2 x D_{coast} = 0.07 g. This selection allows mild deceleration amounts so that the first

order response of the linear forces feels comfortable due to the limited deceleration limit placed on the linear virtual forces.

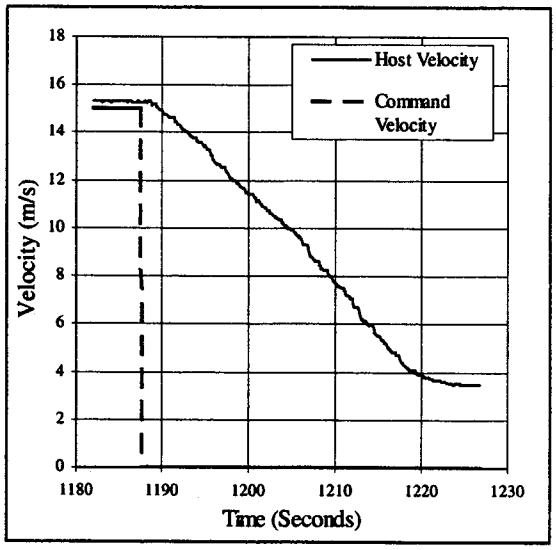

Figure 2.1: Response of velocity controller to a step input of 15 m/s without braking

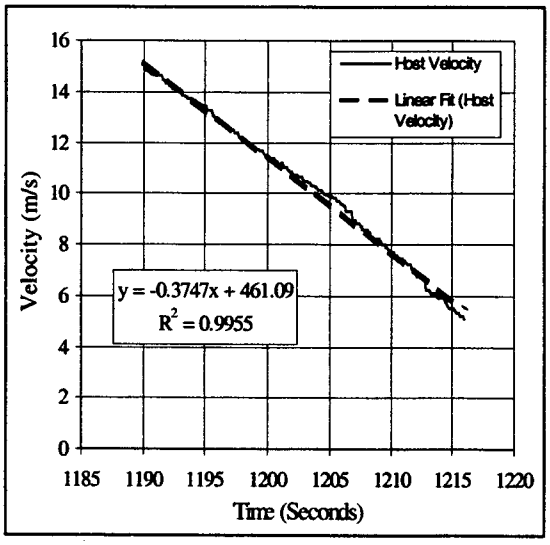

Figure 2.3: Linear fit of deceleration from 15 m/s without braking

The maximum deceleration of the velocity controller was determined in a similar manner. The host was given a velocity command of 20 m/s. Once the host achieved this velocity, the command was stepped down to zero m/s. This time, however, the braking system was active. The velocity controller stopped the truck with the velocity decreasing in a linear fashion as shown in Figure 2.5.

Isolation of the linear region in Figure 2.5 was fitted with a line as shown in Figure 2.6. The maximum velocity controller deceleration with braking was -2.76 m/s^2 (-0.24 g's). This deceleration was designed to be conservative. Kempf at the Transportation Research Center (TRC) tested a Navistar tractor with ABS brakes and reported their result in [32]. Their deceleration experiments were performed from 60 mph on a Navistar tractor with and without a trailer (the trailer had no braking capabilities) on several types of road surfaces. The deceleration

capabilities ranged from 0.64 g's to 0.45 g's. A trailer with brakes would be able to achieve higher deceleration levels. The lowest g-level in all the tests for the tractor without a trailer was 0.56 g's. Clearly, the velocity controller limits the application of the brakes so that the maximum deceleration of the velocity controller is below the capabilities of the Navistar tractor ABS system. The maximum deceleration limit for the virtual bumper ($Dmax_H$) was set at -0.24g's to conservatively match the braking capabilities of the velocity controller. The conservative approach was chosen in the interest of safety and to limit the wear on the brakes and tires.

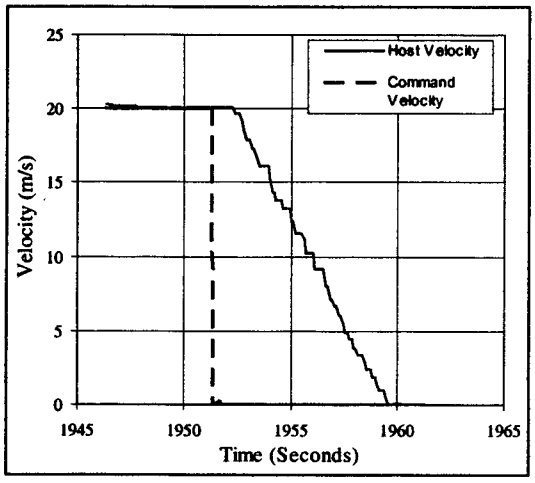

Figure 2.3: Response of velocity controller to a step input of 20 m/s, brakes active

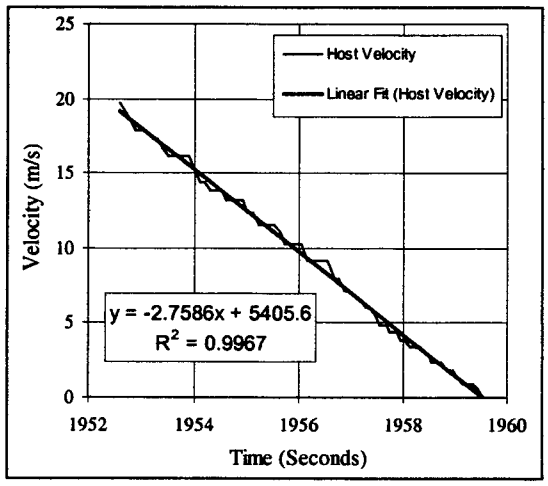

Figure 2.4: Linear fit of deceleration from 20 m/s, brakes active

2.2 Description of Simulation

The virtual bumper was implemented in simulation by Schiller [2]. While this simulation proved that the algorithm was effective in providing a desired response, Schiller did not explore some of the issues involved in a real world implementation. For example, we wanted to determine how robust the algorithm was to sensor noise and latency. Schiller's simulation consisted of thousands of lines of code and was tightly coupled to a three dimensional graphical library which was used to simulate an array of radar units. We do not plan on using this approach for the initial implementation of the virtual bumper. We felt it would be more useful and less time consuming to re-program the algorithm, but still use his truck model.

The lateral truck model used by Schiller was obtained from a previous research effort, which is described in a report written by Lee Alexander et al [33]. Schiller ported the equations to his simulation. The transfer function relating steering input to the yaw angle of the tractor was a fourth order equation. The longitudinal dynamics were modeled by Schiller as a simple second order equation. No experiments were performed to verify their accuracy. Clearly, this is an approximation of the real dynamic response of the truck. The virtual bumper is a higher level system that relies upon the subsystems to contain the accurate dynamic models of the steering and throttle/braking system. Since these components were already installed and working on the Navistar, we did not invest a great deal of time investigating the accuracy or validity of the dynamics models. We used the simulation to find out how the virtual bumper gain settings and sensor noise affected the behavior of the longitudinal controller.

The approach taken to programming the algorithm in simulation was to take an implementation based approach, that is, to minimize the software changes needed to port the simulation to the real hardware aboard the tractor. To accomplish this, the simulation was programmed in VxSim, which is a simulation of the real time operating system VxWorks. VxWorks is the operating system used on the SAFETRUCK. VxSim has function calls identical to those for VxWorks and code written in VxSim is virtually 100% portable to VxWorks. The simulation was also programmed using the same Inter-Process Communication (IPC) Application Programming Interface (API) and data structures for software already written and tested aboard the tractor. This made the port between simulation and the real system seamless.

Figure 2.5 shows the computer architecture of the system. The simulation was programmed on a UNIX workstation running the SunOS. The simulated VxWorks operating system, VxSim, resides on the UNIX workstation as a user process. The computer system aboard the Navistar tractor resides on a VME bus and consists of three Motorola MVME-147 processors. VxWorks is the operating system (OS) on this computer system. The velocity controller was programmed

on a x86 PC laptop running the Real Time Kernel (RTK) OS [31]. The VME communicates with the velocity controller via serial communications.

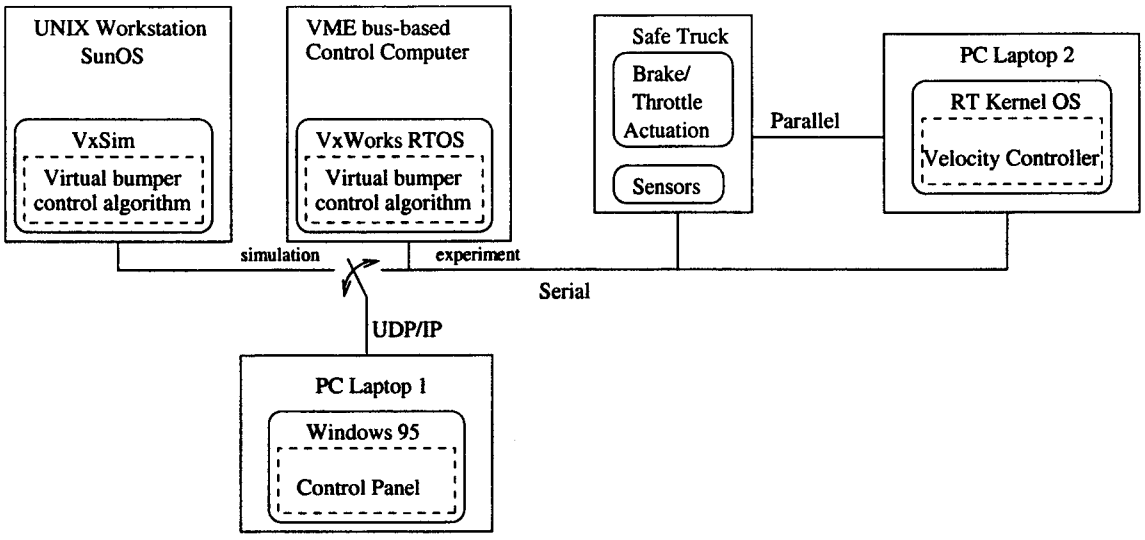

Figure 2.5: Computer Architecture of Computer Simulation and Real System

To facilitate the analysis and control of each experiment in real time, a control panel was programmed in the C++ programming language. The scripting language Tcl/Tk was used for the Graphical User Interface (GUI). The control panel receives sensor information from the virtual bumper via UDP/IP internet communication. We implemented a local Ethernet LAN aboard the truck so that the host, target and control panel can communicate.

The goal behind the design of the control panel was to graphically display the location of the target(s) and to be able to adjust controller gains in real time. To that end, the control panel draws a range vs. range rate phase plot and a birds eye view of the driving scene (see Figure 2.6) in real time. The top view includes the host, all personal space boundaries, targets and the virtual forces acting on the host. The range vs. range rate window provides dynamic vehicle interaction data. Although this document has been printed in black and white, the actual display panel is in color. In the range vs. range rate plot, the R/Rdot point is black while the linear personal space

range is in yellow and the non-linear personal space range is in red. The personal space boundaries in the top view have the same color code.

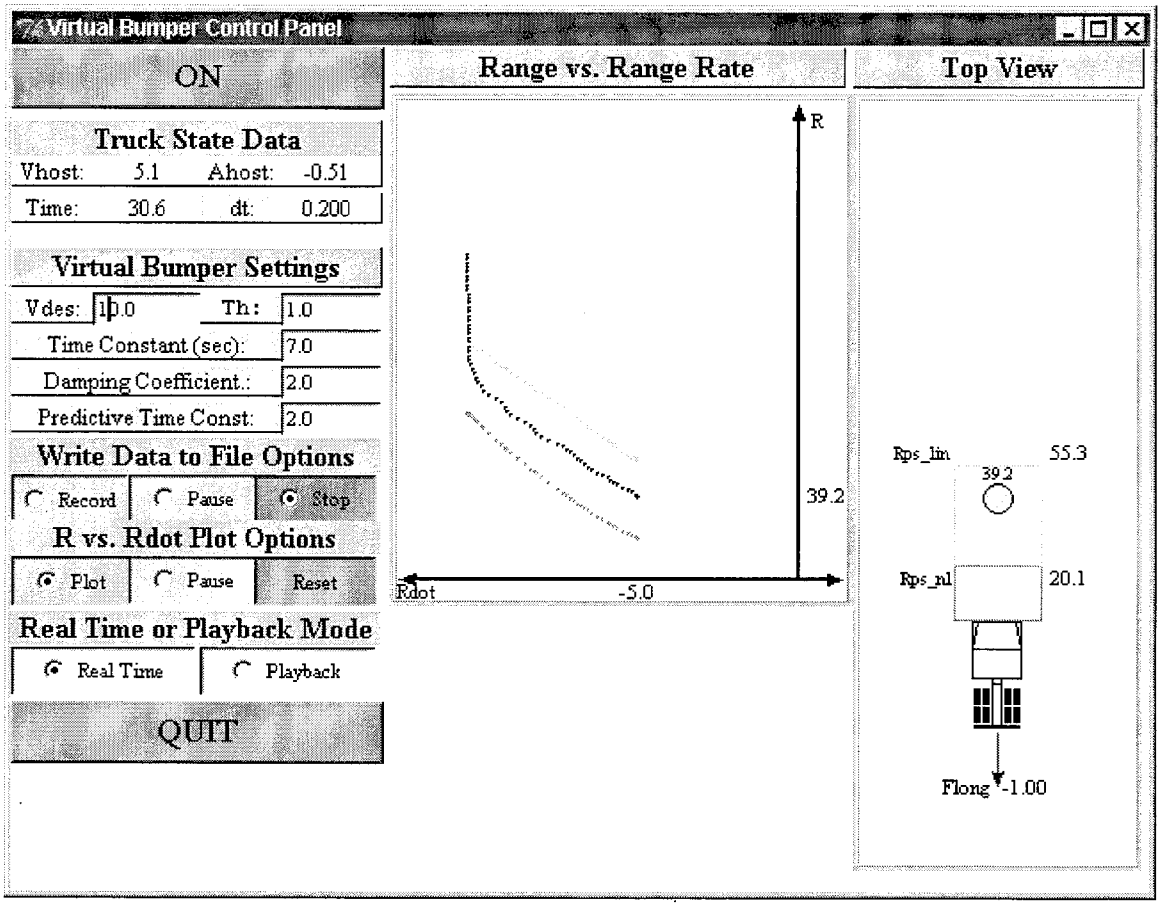

Figure 2.6: Screen shot of virtual bumper control panel. All units are SI.

The "Virtual Bumper Settings" area allows us to modify the experiment by choosing the desired host velocity, headway time, and controller gains. These values can be changed "on the fly" without recompilation of the source code or modifying initialization files.

2.3 Effect of Gains on System Response

To determine how each gain affects the behavior of the controller, we ran the simulation and adjusted each gain individually. This provided us with a feel for each gain's effect on the behavior of the host vehicle so that tuning the actual system would be more intuitive.

2.3.1 Effect of the Headway Time Constant, τ

The dynamics of the headway controller are determined by the characteristic equation. By design, we have forced the solution of the characteristic equation to be stable and overdamped. The pole nearest the origin is termed the dominant pole and has the greatest effect on the system behavior. We place the dominant pole through the parameter τ ($1/p_1$), the headway controller time constant.

This parameter has the greatest effect on the overall response of the controller. It determines how 'reactive' the controller will be. Fancher [34] chose a time constant of 12 seconds because his system did not incorporate braking. Our system includes braking and throttle so we can select a smaller time constant because our system can produce higher deceleration levels.

For the remainder of this thesis, the range vs. range rate phase plot and a time domain plot of the major variables will be used to characterize the response of the headway controller for each experimental/simulated run. We have previously described the range vs. range rate phase plot and its importance in graphically presenting the behavior of the host-target vehicle headway relationship. The time domain plots incorporate; the range of the target, the desired range (Rh), the target and host velocity, the acceleration and virtual force. The forces were normalized by dividing the virtual force by the value of the mass of the truck. This was done so that the same scale could be used for all the time plots.

To determine the specific effects of the time constant parameter, we varied it in simulation. We present the results for τ = 9.0 seconds and τ = 5.0 seconds (Figure 2.7- Figure 2.10). First, we will define the curves that appear on every subsequent range vs. range rate phase plot. The 'switching line' is the personal space boundary and under it the linear and nonlinear forces are applied to the host. The 'R_{des}' line shows the desired $(R, \dot{R})$ trajectory as defined by Equation 2.1. The 'D_{psnl}' curve defines the boundary between the linear and nonlinear force regions. Finally, the curve '$Dmax_h$' defines the maximum deceleration capabilities of the velocity controller.

The range vs. range rate phase plots (Figure 2.7- Figure 2.10) show that the higher τ value increases the range at which the controller is activated (slope of the switching line, which is the maximum range of the linear personal space). The desired response is determined by the following;

$$R_{des} = -\tau\dot{R} + R_H. \qquad (2.1)$$

It is clear that the desired response line has a slope of τ ($1/p_1$). The larger τ increases the slope of the desired response, which activates the headway controller at larger range.

The time domain plots reveal that the virtual longitudinal forces and host acceleration are higher for $\tau = 5.0$ than $\tau = 9.0$. This is expected because the control action starts at a lower range for the $\tau = 5.0$ simulation and the lower slope of the desired response line indicates that $\dot{R}$ changes at a faster rate, which requires a higher deceleration. Furthermore, notice that the time for the range to settle to R_H is lower for the $\tau = 5.0$ simulation than the $\tau = 7.0$.

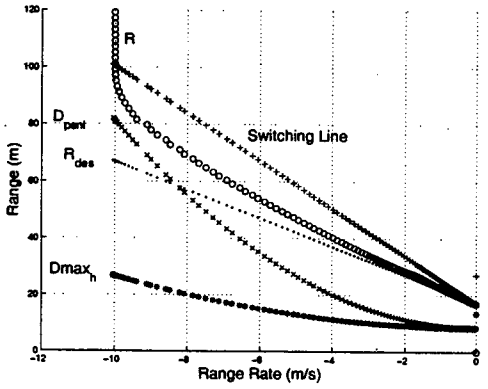

Figure 2.7: R vs. $\dot{R}$; Vhost = 25 m/s, Vtarget = 15 m/s, τ = 5 sec

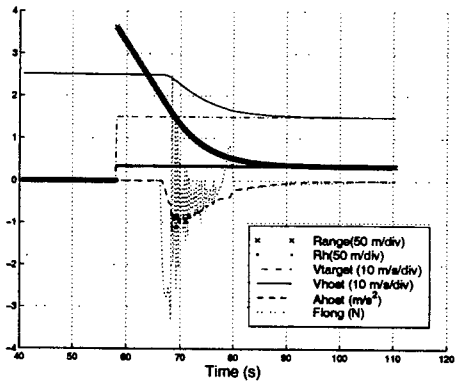

Figure 2.8: Time domain plot; Vhost = 25 m/s, Vtarget = 15 m/s, τ = 5 sec

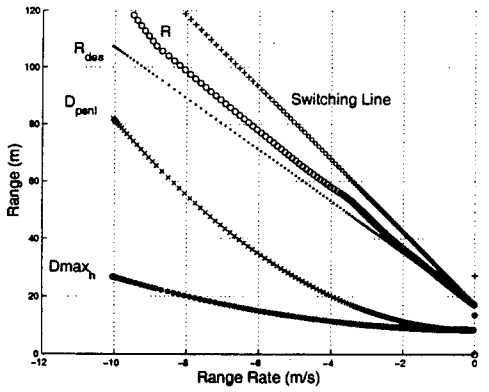

Figure 2.9: R vs. $\dot{R}$; Vhost = 25 m/s, Vtarget = 15 m/s, τ = 9 sec

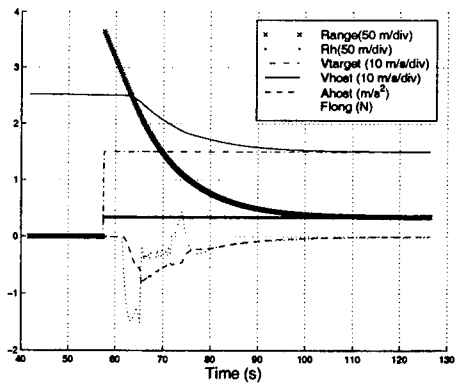

Figure 2.10: Time domain plot; Vhost = 25 m/s, Vtarget = 15 m/s, τ = 9 sec

The general effect of τ, as expected, is to change the controller response time. A higher time constant produces lower forces and increases response time. Conversely, a low time constant produces higher forces and a faster response. From the analysis of a number of simulation runs, we chose a time constant of 7.0 seconds, lower than the 12 second time constant selected by Fancher.

2.3.2 Effect of the Damping Ratio, ζ

While τ determines the general response time of the longitudinal controller, the damping ratio determines the location of the second pole. Once τ and ζ are chosen, one can solve for ω_n, p_2, b and k.

The choice of ζ and its effect on the controller response can be demonstrated by varying ζ in the simulation while holding all other parameters constant (Figure 2.11 - Figure 2.14). The range vs. range rate phase plots for different values of ζ reveal several interesting facts. First of all, notice that the switching line (linear personal space boundary) decreases with increasing ζ. This is due to the fact that linear personal space range is a function of k and b. Secondly, the parabolic section of the trajectory is reduced as ζ increases. This is due to the second order effect caused

by the second pole, p_2. The larger ζ becomes, the further away the second pole is located from the origin of the s domain phase plane and the less effect it has on the response. The system response becomes more linear (first order) and more closely follows the desired response line determined by the location of the dominant pole.

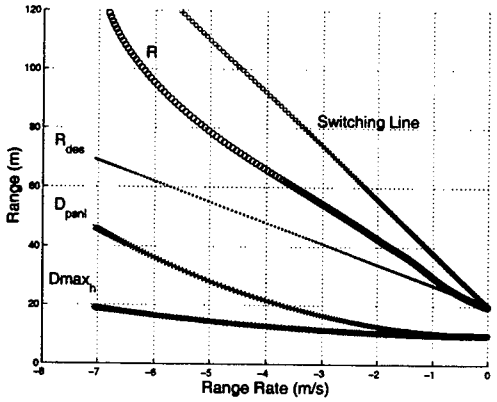

Figure 2.11: R vs. $\dot{R}$ plot with Vhost = 25 m/s, Vtarget = 18 m/s, $\tau = 7.0$, $\zeta = 1.0$

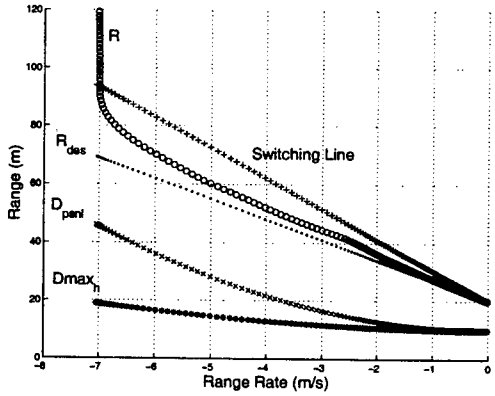

Figure 2.12: R vs. $\dot{R}$ plot with Vhost = 25 m/s, Vtarget = 18 m/s, $\tau = 7.0$, $\zeta = 2.0$

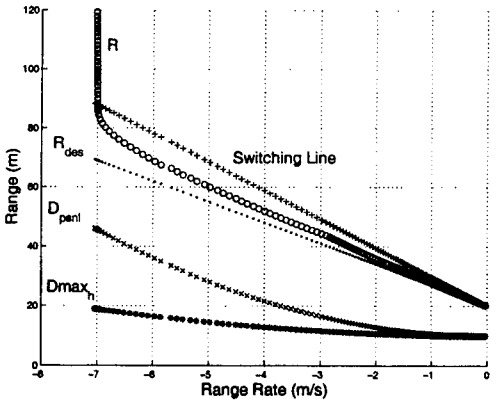

Figure 2.13: R vs. $\dot{R}$ plot with Vhost = 25 m/s, Vtarget = 18 m/s, $\tau = 7.0$, $\zeta = 3.0$

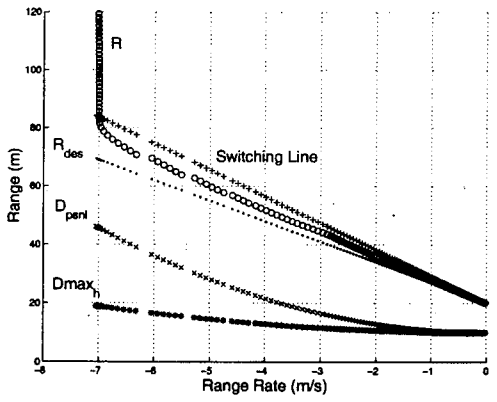

Figure 2.14: R vs. $\dot{R}$ plot with Vhost = 25 m/s, Vtarget = 18 m/s, $\tau = 7.0$, $\zeta = 4.0$

The time domain plots (not shown) indicate that the longitudinal forces increase with increasing ζ. A higher longitudinal force is desired as long as the response is smooth and has little overshoot. It would be optimal to have ζ as high as possible without the large forces that may trigger oscillations or even instability in the actual system. For this simulation, $\zeta = 2.0$ gives a desired response with reasonably low virtual forces.

2.3.3 Effect of the Predictive Time Constant, T

Once τ and ζ have been selected, the predictive time constant (T) is set. The approach is to increase T from zero until the headway response is smooth with no overshoot. T affects the personal space boundary as well as the linear longitudinal force applied to the host. To determine how T affects the response of the headway controller, we held τ and ζ constant and varied T from 0.0 to 4.0 seconds. The result for the host traveling at 25.0 m/s and the target at 18 m/s is shown in Figure 2.15 - Figure 2.18.

It is clear that without the predictive time constant (T = 0.0 sec), the longitudinal controller is oscillatory due to the tracking error of the velocity controller. The headway oscillates around the desired headway as indicated by the circular pattern of the range-range rate trajectory.

Increasing the predictive time constant to T = 1.0 sec eliminated the oscillation and overshoot but the response does not follow a straight-line trajectory to the desired headway. When T is increased to 2.0 sec, the response is mostly linear without oscillation and overshoot. When T is increased even further to 3.0 sec, the response is even more linear. Notice that the linear personal space boundary increases with increasing T. Therefore, T should be chosen as the minimum value that will produce a reasonably linear response and eliminate oscillation and overshoot as well as minimize the maximum sensing requirements of the range sensors. A predictive time constant of 2.0 sec produces the best results given these criteria.

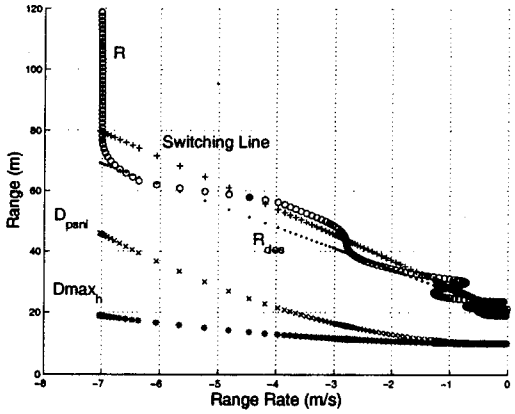

Figure 2.15: R vs. $\dot{R}$ plot with Vhost = 25 m/s, Vtar = 18m/s, τ = 7.0 sec, ζ = 7.0, T = 0.0 sec

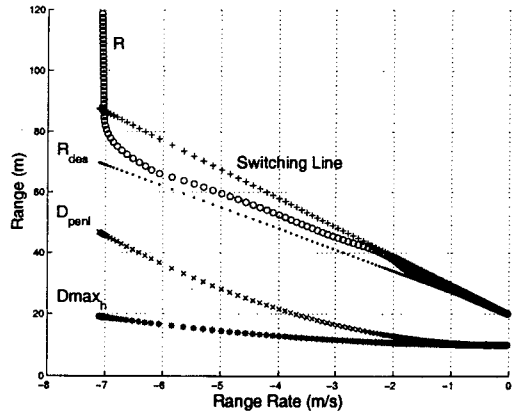

Figure 2.16: R vs. $\dot{R}$ plot with Vhost = 25 m/s, Vtar = 18m/s, τ = 7.0 sec, ζ = 7.0, T = 1.0 sec

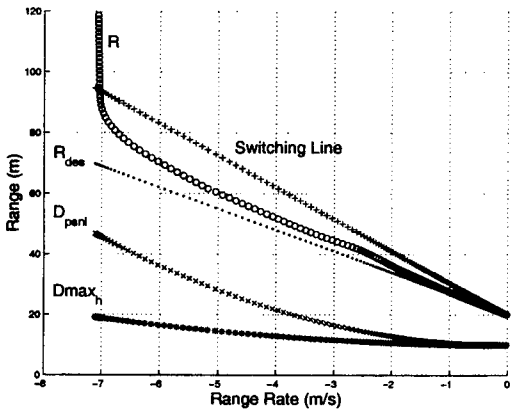

Figure 2.17: R vs. $\dot{R}$ plot with Vhost = 25 m/s, Vtar = 18m/s, τ = 7.0 sec, ζ = 7.0, T = 2.0 sec

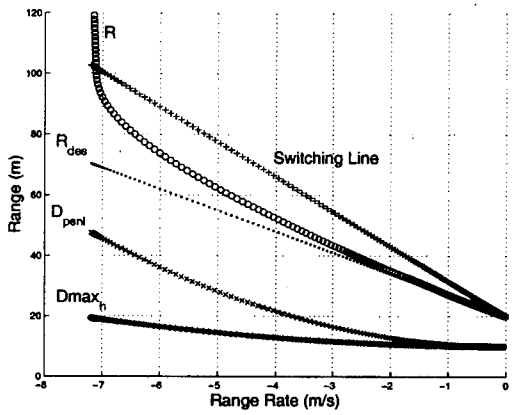

Figure 2.18: R vs. $\dot{R}$ plot with Vhost = 25 m/s, Vtar = 18m/s, τ = 7.0 sec, ζ = 7.0, T = 3.0 sec

2.4 Effect of Sensor Noise and Latency

Thus far, the simulation has been run with the assumption of noiseless sensors with no latency. Since we will implement the controller on a real truck, we were interested to see how robust the controller is to sensor noise and latency. This information is useful in determining what filters

are needed and what software synchronization scheme should be used. In this section we present the results of the simulated controller with sensor noise and sensor latency.

2.4.1 Effect of Radar Noise

In previous chapters we detailed the results of an evaluation of the EVT-200 radar. We planned on using this sensor for the virtual bumper longitudinal controller experiments. The sensor was generally reliable, but exhibited occasional dropouts (loss of target) and increased error at large range. We used this observation to develop a somewhat realistic noise signal for the simulated radar.

We modeled the radar noise as a high frequency sine wave that has an amplitude that is a linear function of the range. This means that larger range values have a greater noise component. This is similar to the behavior exhibited by the EVT-200 evaluation. We also added randomly occurring dropouts. We define a dropout as a range value of zero after the target has previously been detected and the headway is clearly within the maximum sensing range of the radar.

For all the simulations in §2.3, the host velocity was set to 25 m/s and the target was traveling at 18 m/s. The gains were held constant and set to the values we determined to be optimal in the previous sections: τ = 7.0 sec, ζ = 2.0 and T = 2.0 sec. The results in Figure 2.19 and Figure 2.20 show that the longitudinal controller is robust to range sensor noise of this type. The range vs. range rate phase plot shows an oscillatory trajectory. This normally would not be desirable, but consider that this is caused by the error in the range measurement. The host velocity in the time domain plot shows a smooth reduction to the target velocity. The dropouts also had little effect on the desired response. The dropouts are clearly shown in Figure 2.20 by the target velocity plot. When a dropout occurs, the target velocity jumps to zero because no target is detected. This dramatic change caused the longitudinal force to change rapidly, but the longitudinal admittance and the mass of the host vehicle prohibit instantaneous velocity changes.

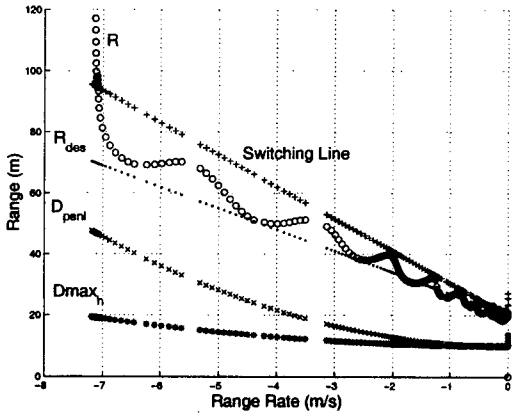

Figure 2.19: R vs. $\dot{R}$ plot with sinusoidal radar noise with dropouts

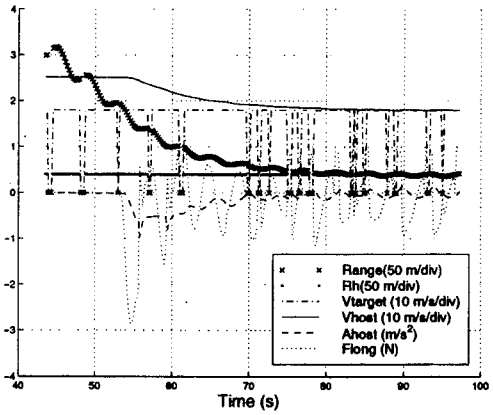

Figure 2.20: Time domain plot with sinusoidal radar noise with dropout

2.4.2 Effect of DGPS Noise

The host vehicle state is determined from DGPS sensor data. The velocity and acceleration are obtained from the time derivative and the double time derivative of successive location data, respectively. Thus, any error in the DGPS signal will be translated (and amplified) into host velocity and acceleration error. The linear force equation uses this information; thus, DGPS noise effects the forces imparted upon the host vehicle.

To determine how sensitive the force equation is to DGPS sensor noise/error, we modeled the sensor with Gaussian noise. This does not accurately model actual DGPS error. Bajikar et al [35] measured the dynamic error to be 25 cm in the longitudinal direction and 5 cm in the lateral direction. Clearly, the Gaussian model of noise does not fit this conclusion. Gaussian noise will produce the same error (on average) in both the longitudinal and lateral directions. Regardless, we are only concerned with the error in velocity and acceleration and can ignore errors in the lateral direction. As long as the average noise is greater than 25 cm, the noise will approximately simulate the noise in DGPS.

The result of a simulation with DGPS noise is shown in Figure 2.21 and Figure 2.22. The range vs. range rate plot shows that the host vehicle followed an acceptable trajectory to the desired headway. The time domain figure reveals that the forces increased dramatically with DGPS noise. While this force increase had little overall effect on the simulated host dynamics, large forces may affect the real system if the modeled and actual truck dynamics differ significantly.

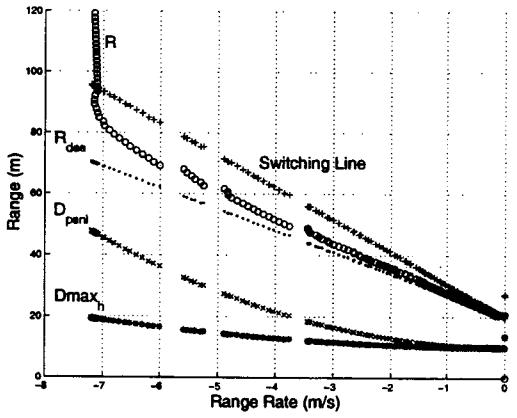

Figure 2.21: R vs. $\dot{R}$ plot with Gaussian DGPS noise

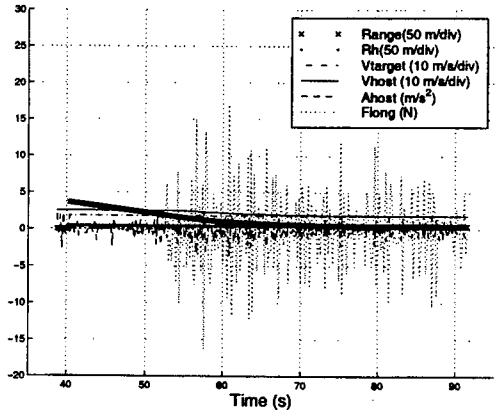

Figure 2.22: Time domain plot with Gaussian DGPS noise

2.4.3 Effect of Radar Latency

Timing in a real time control system is critical. Latency due to CPU overload, I/0 and sensor processing can cause unexpected behavior and even instability in the control system. Data collection schemes can affect the latency of some sensor data. It is important, therefore, to determine which sensor measured data causes the greatest undesirable effect in the control algorithm.

We tested the sensitivity of the virtual bumper longitudinal controller to delays in receiving radar data through the simulation. This was done by delaying the simulated radar data by m*0.1 seconds, where m is a non-negative integer. All other sensor data was acquired with zero latency. The result of the host traveling at 10 m/s towards a stationary target is shown in Figure

2.23 and Figure 2.24 (all gains as before). The simulation showed that the algorithm is not very sensitive to relatively large delays in receiving radar data. This means that it is not critical to synchronize the sensor data acquisition on radar.

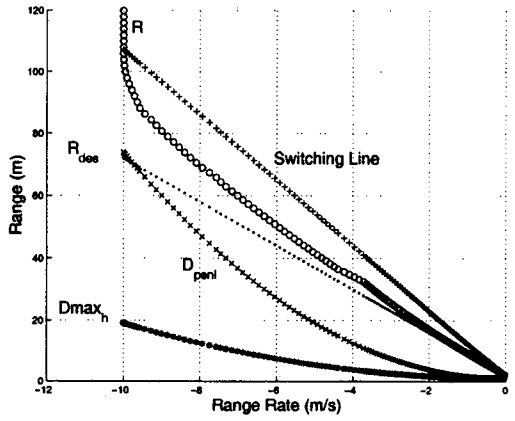

Figure 2.23: R vs. $\dot{R}$ plot with no radar delay

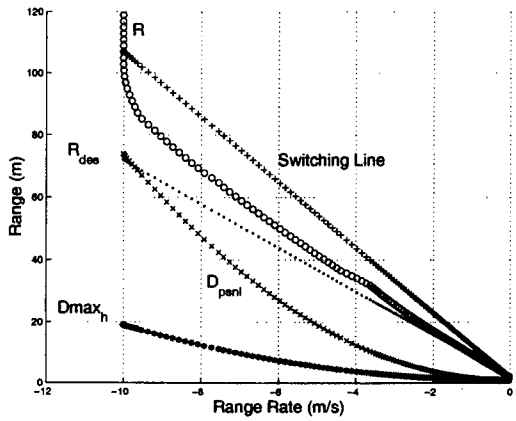

Figure 2.24: R vs. $\dot{R}$ plot with radar delay = 0.3 sec

2.4.4 Effect of DGPS Latency

Since the DGPS data acquisition rate is half that of radar, or 5Hz, it was desired to synchronize data collection based on radar. This would provide a bandwidth of 10 Hz for the virtual bumper. Unfortunately, the algorithm proved more sensitive to DGPS (host velocity, acceleration) latency than radar latency.

The simulation was used to verify this effect. Simulated DGPS data was delayed by increments of 0.2 seconds. The gains were the same as previous simulations and the host traveled at 10 m/s towards a stationary target. The result is shown in Figure 2.25 - Figure 2.28. This simulation clearly shows that the algorithm is sensitive to DGPS latency and even becomes unstable for large latencies. This sensitivity was the basis for the decision to synchronize sensor data collection on DGPS, not radar. The decreased bandwidth (10 Hz to 5 Hz) is unfortunate, but the improved stability by minimizing DGPS latency makes this compromise acceptable. By

synchronizing sensor data collection on DGPS, we will be able to limit the latency to data transmission delays, which are usually negligible.

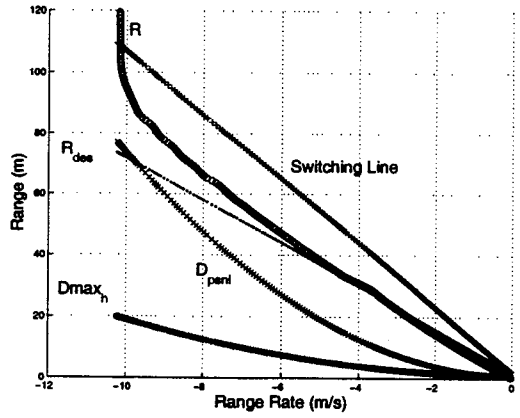

Figure 2.25: R vs. $\dot{R}$ plot with no DGPS latency

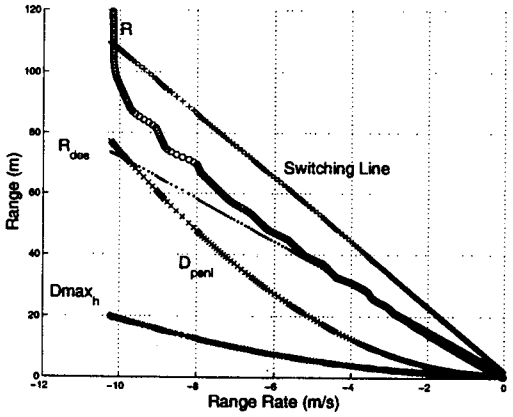

Figure 2.26: R vs. $\dot{R}$ plot with DGPS latency = 0.2 sec

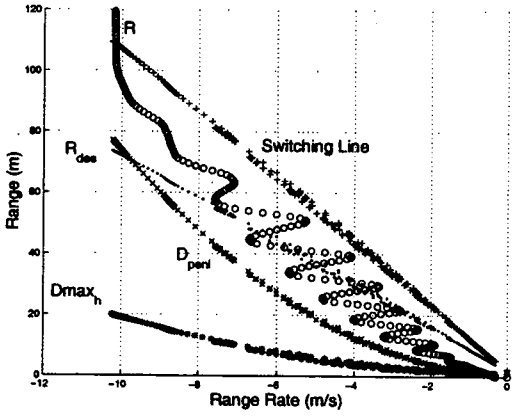

Figure 2.27: R vs. $\dot{R}$ plot with DGPS latency = 0.4 sec

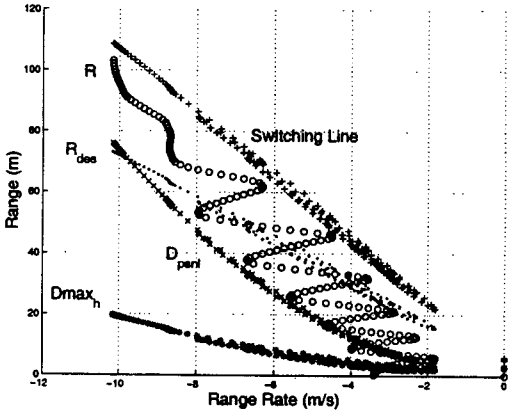

Figure 2.28: R vs. $\dot{R}$ plot with DGPS latency = 0.6 sec

2.5 Simulation Results

We determined from the simulation how each parameter of the longitudinal controller affects the host/target headway relationship. The controller time constant τ determines the time constant of

the control action. Higher values of τ result in a longer time for the host to achieve the desired headway. Smaller values have the opposite effect.

The damping ratio, ζ, determined how closely the host follows a linear trajectory on the range vs. range rate phase diagram. Higher values of ζ produce higher force and more linear response. Lower values of ζ have reduced forces and a more contoured trajectory.

The predictive time constant T is tuned to match the velocity controller time constant. Increasing T reduces oscillations in host velocity and increases the personal space. A value for T should be chosen as low as possible to reduce the personal space, but high enough to eliminate host velocity oscillations due to the tracking error of the velocity controller.

We also evaluated the effect of sensor noise. Radar sensor noise was modeled with a high frequency component that had an amplitude that scaled with range. Signal dropouts were also modeled. The longitudinal controller was able to successfully reduce the headway to the desired value. The DGPS noise was modeled as Gaussian. Since host velocity and acceleration are derived from this sensor, the knowledge of the effect of DGPS noise was critical. The host arrived at the desired headway, but virtual forces were dramatically increased.

To determine the effect of sensor latency on the controller's ability to control the headway distance to the target, we used the simulation to delay the radar and DGPS sensor data. The effect of radar sensor latency was negligible. However, DGPS showed more sensitivity to sensor latency, and instability can occur if the delay is large. It was decided that software data collection synchronization routines would be synchronized by DGPS, minimizing its latency on the virtual force calculation.

Chapter 3

Experimental Results

3. Experimental Results

In this chapter we describe the experiments performed with the virtual bumper longitudinal controller. We present the results of experiments that were designed to test the controller in various real world driving situations. An analysis of each experiment is given as well as a discussion of the strengths and weaknesses of the controller for each driving scenario.

3.1 Experiment Design

The main philosophy behind the longitudinal controller experiments is to simulate real highway driving scenarios. Thus, we chose several driving situations that we felt would tax the longitudinal controller so that the algorithm is proven robust and effective for both normal and emergency driving.

We divided the driving scenarios into four main categories (Figure 3.1);

1) Adaptive Cruise Control (ACC)
2) Major traffic slow down
3) Critical stop
4) Miscellaneous driving scenarios

The ACC functionality of the longitudinal controller is essential in highway driving. ACC incorporates a standard cruise controller with a headway controller. When the host approaches a vehicle moving slower than its user set cruising speed, the host vehicle slows down by releasing the throttle to match the velocity of the target vehicle. In the major traffic slow down scenario, the target vehicle is traveling significantly slower than the host. The host must apply braking to reduce its velocity to match the target. In the critical stop scenario, the target vehicle is at rest. This simulates such situations as approaching a severe traffic jam or a stalled/abandoned vehicle in front of the host. This will test the longitudinal controller's ability to stop the host with a safe final headway. The miscellaneous driving scenarios represent other less frequent driving situations. Some examples include a target vehicle cutting in front of the host and stop and go traffic.

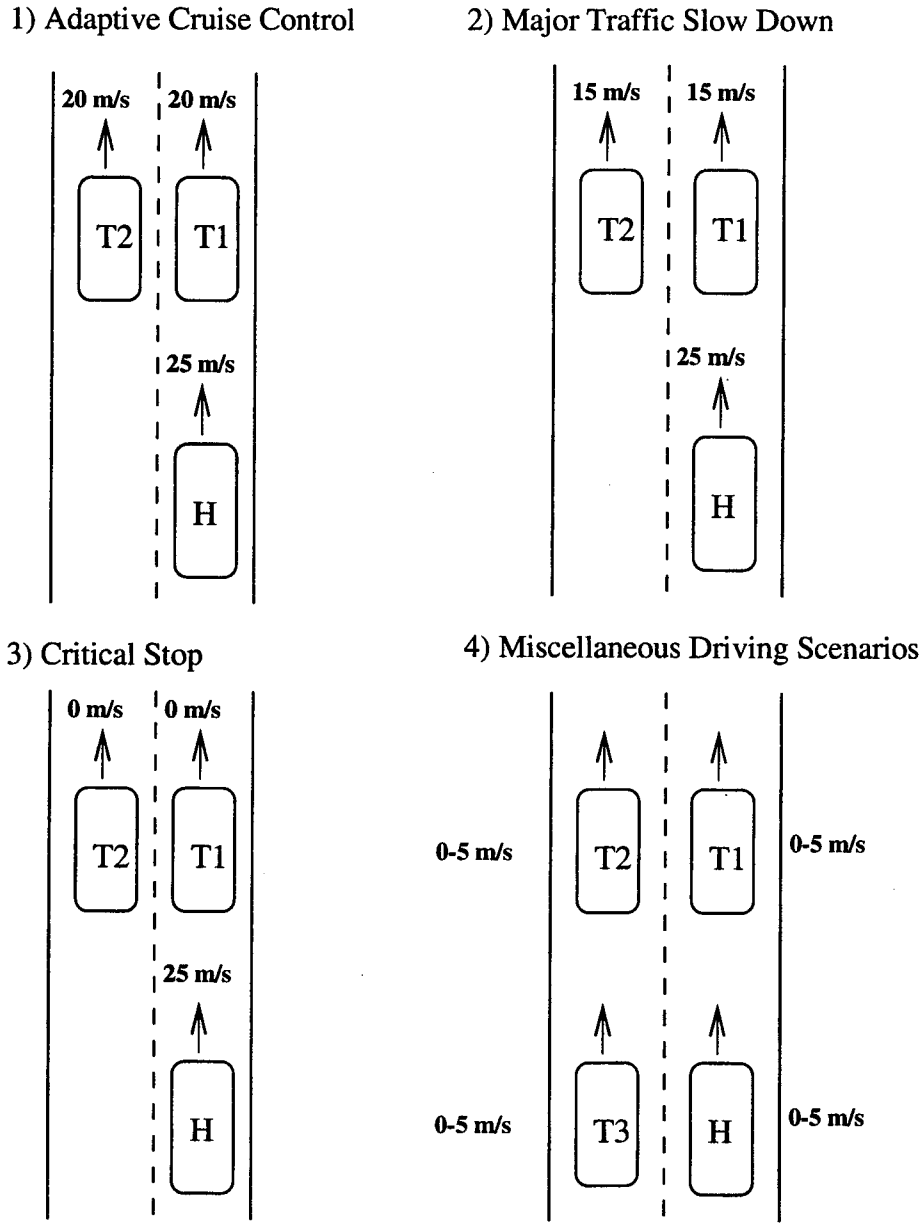

Figure 3.1: Driving scenarios to be tested

3.2 Algorithm Modifications

Simulations are useful to verify the effectiveness of control algorithms both assuming ideal conditions and conditions involving structured noise and latency. The real world, however, has

issues that are difficult to model in simulation and thus unexpected behavior of the control system commonly occurs. This was the case for the longitudinal controller.

Unexpected behavior was discovered during the initial ACC experiments and was caused by the definition of linear personal space presented in Chapter 1. The linear personal space (R_{ps_lin}) was defined by a line that intersects the range rate axis in the range vs. range rate phase plot. This point of intersection is the desired headway, R_H, when steady state conditions are reached. Sensor noise and latency, transmission gear shifts and tracking errors in the velocity controller can cause slight deviations from the desired state. As the personal space boundary was defined in the previous chapter, any deviation to the right (positive $\dot{R}$) of the desired headway in the range vs. range rate phase diagram results in the target vehicle exiting the personal space boundary (see Figure 1.10). When this happens, the longitudinal controller is in cruise control mode and the cruise control forces push the host in an attempt to accelerate it to the desired traveling speed. This force is applied until the target enters the personal space region. Momentum of the host as well as the inherent delay of the velocity controller cause the host to accelerate towards the target even after it is within the personal space and the forces becomes negative. The negative forces become large as the target falls below the desired headway range until the host finally decelerates (the velocity controller reacts to the velocity command). It was observed that the host slows down until $\dot{R}$ becomes positive. The target then falls out of the personal space and the cycle repeats itself.

The resultant oscillation of headway is neither comfortable for the driver nor does it imitate how human drivers follow a vehicle. Thus, we increased the personal space boundary by a constant buffer distance given by the following equation.

$$R_{PS\_nl} = T\dot{R} + R_H + R_{buff} \tag{3.1}$$

The additional term, R_{buff}, allows the target vehicle to slightly deviate from the desired headway (with $\dot{R}$ positive) without invoking the cruise control forces. In effect, the added term raises the

switching line in the range vs. range rate phase plot so that R_H is no longer at the intersection of the personal space boundary. The linear forces are smaller than the cruise control forces (especially near R_H) and more gently accelerate the host until it again approaches the desired headway. This alteration to the personal space effectively prevents the large headway oscillations caused by the target leaving the personal space at the slightest deviation from steady state. This provided the controller with greater disturbance rejection at steady state.

Experiments determined that a R_{buff} of 6.0 m was sufficient to prevent the target from leaving the personal space due to small deviations of range. The side effect of this added personal space is that targets first entering the personal space boundary will receive a small positive force because the equation is no longer based on driving the linear force to zero (see §1.4.2). Since this buffer distance is small compared to the personal space boundary, the target is only in this region for a brief period of time. Furthermore, the target state is close to the $F_{lin} = 0$ line (the previous definition of the personal space boundary) and the resultant positive forces are small. The effect caused by positive forces resulting from the increase in the personal space was negligible in the experiments performed.

3.3 Range Sensing Issues

The goal of this thesis is to validate the control laws of the virtual bumper longitudinal algorithm. Therefore, a greater emphasis was placed on the algorithm itself and not on the sensing modality. The sensors used for testing the longitudinal controller were adequate for this purpose. The algorithm is dependent upon knowledge of the surrounding driving environment. We will document what sensor characteristics are important for the algorithm to function properly, but we did not spend a great deal of time designing the optimum sensing configuration.

An Eaton VORAD EVT-200 radar and an Optech Sentinal 100 laser range finder were used to sense the target vehicle in the front personal space region. Both sensors are unidirectional, so we only tested the system on straight portions of the test track. Modern range sensors provide information for the direction (azimuth) to the sensed target making them more effective on

curved roads. The virtual bumper can be easily adopted to curved roads by 'warping' the personal space regions to match the curvature of the roads. Since we currently do not have the sensing capability for handling obstacles or vehicles located ahead on a curved road, the algorithm was only tested on straight roads.

The Eaton VORAD radar is a Doppler sensor making it ideal for detection of emergency driving situations (high relative velocities). It performs poorly, however, during ACC maneuvers because the relative velocity between vehicles is small. It was for this reason that this radar was used for the critical stop experiments only. The radar has a four degree azimuth cone angle, but the location of the target is not provided, only range and range rate are measured. This causes false target signals at large range because objects on the side of the road are detected. This limits the effective maximum range of the EVT-200 radar.

The Optech laser range finder is a time of flight sensor. It measures the range to target regardless of the relative velocity between objects. Hence, this sensor performed well in ACC driving scenarios. The negative aspect of this sensor is that it does not directly measure the range rate to the target and its ability to detect vehicle under heavy snowfall has not been tested. The range rate to target was calculated using successive range measurements and the measured time between them. Unfortunately, range errors caused the range rate calculation to be noisy at long range. At closer range, the range measurements had higher accuracy, which reduced the noise in the range rate calculation. It will be shown that the longitudinal controller was more robust to sensor noise than to false targets. The laser range finder has a very narrow beam so that false targets at long range were not an issue on straight portions of the track. It was for this reason, along with its insensitivity to relative velocity, that the Optech range finder was used in the majority of our experiments.

3.4 Controller Gains

The spring stiffness and damping coefficient values were first set to match the best results achieved in simulation. In the simulation, the time constant of the controller (τ) was chosen to be 7.0 seconds, the damping ratio (ζ) was 2.0, and the predictive time constant (T) was 2.0

seconds. This set of gains demonstrated the most desirable in response in simulation given the design criteria.

After implementing the virtual bumper longitudinal controller aboard the Navistar tractor, it was immediately observed that the virtual forces were much higher than those calculated in the simulation, given the same gain values. The result was an oscillatory behavior in velocity and headway.

The large forces were attributed to DGPS sensor noise and differences in the simulated velocity controller and the actual velocity controller. The simulations described in our previous chapter predicted that the DGPS noise would cause larger virtual forces. What was not as obvious is the difference between the velocity controller model and the actual velocity controller. The velocity controller was able track velocity commands (especially deceleration using the brakes) more closely than the simulated truck model. Thus, large drops in the velocity commands, caused by the increased forces, that were not observed in the simulation were achieved in the real system. This caused the truck to slow down more than in the simulation, which moved the target state $(R, \dot{R})$ to a position above the desired trajectory. The forces then became positive in order to reduce the headway and decrease $\dot{R}$ to the desired state. This caused the host to accelerate toward the target until the forces became negative, which would slow the host and force the target state to the desired trajectory. The velocity controller then applied the brakes to try and achieve the velocity command. It again slowed the vehicle down more than expected. This cycle repeated itself until the desired headway was reached.

The aforementioned behavior was not desirable and not part of our stated design objectives. It is important that the deceleration be non-oscillatory and with little or no overshoot so that the ride is comfortable. Thus, the damping ratio was decreased to 1.0. The controller and predictive time constants were held the same as in the simulation, 7.0 and 2.0 seconds, respectively. Experiments performed with these gain settings showed reduced forces and over-damped behavior.

3.4 Adaptive Cruise Control Experiments

The controller was first tested on the low volume road of the Minnesota Road Research Facility (Mn/ROAD) pavement research facility. The closed loop test track has two straight portions that are long enough to test the ACC functionality at lower speeds, but the track length is insufficient to test the controller at highway speeds. Thus, high speed experiments were performed on the main line of Mn/ROAD, which is a section of highway where traffic from the west bound lane of interstate 94 is diverted for several miles and then returned. This gave us an ample stretch of road on which to test the algorithm at highway speeds and to observe its behavior as it maintained a steady state headway.

The result of a low speed experiment in which the target vehicle was traveling at 9 m/s (32.4 kph) is shown in Figure 3.2 through Figure 3.4. The headway time was set to 1.0 second in this experiment and the host's desired cruising speed was 18 m/s (64.8 kph). Notice in both plots that the host velocity increases until the target state is below the switching line. This is due to the limited length of track on the low volume road, which made it difficult to achieve the desired speed before catching up to the slower moving target vehicle. The range vs. range rate phase plot shows that the target remained in the linear force region. The R, $\dot{R}$ trajectory is linear and converges upon the desired response line (R_{des}) until the range rate becomes small. There is a small velocity oscillation as the range decreases to the desired state. The time domain plot shows that the range decreases to R_h (~11 m) with no overshoot. It is even clearer in this figure that the host velocity has a small oscillation as it approaches the velocity of the target vehicle.

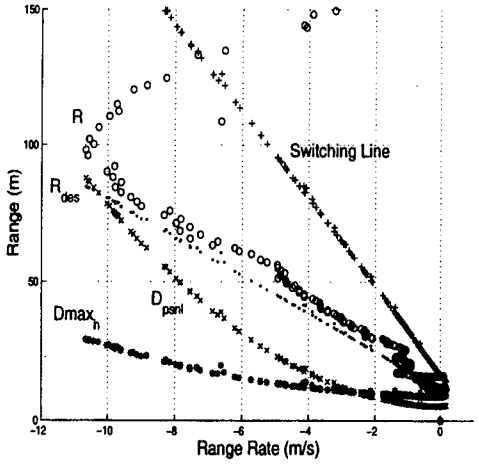

Figure 3.2: R vs. $\dot{R}$ plot, low speed ACC functionality with Vtar = 9 m/s

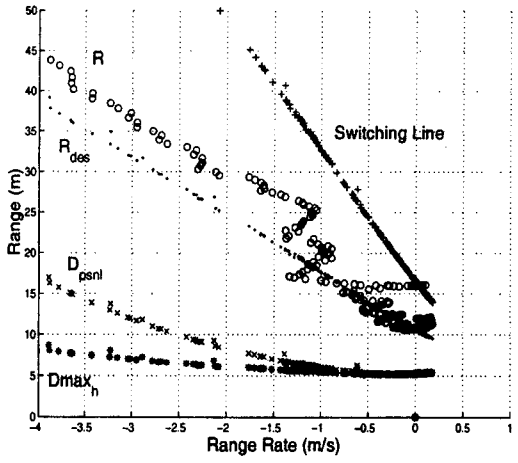

Figure 3.3: R vs. $\dot{R}$ plot, low speed ACC, scaled to show the end of the maneuver

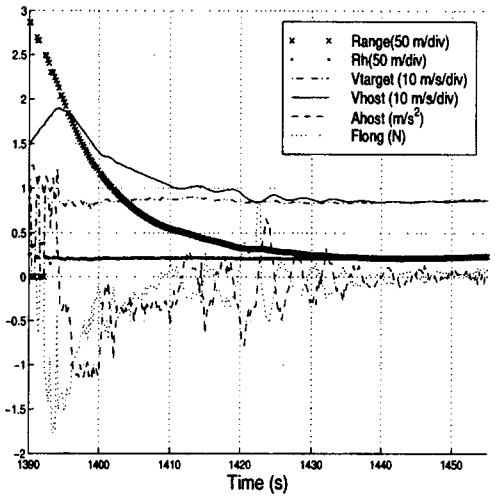

Figure 3.4: Time domain plot, low speed ACC functionality with Vtar = 9 m/s

The small velocity oscillation was not felt by the passengers during the experiment; the deceleration felt smooth. A look at the acceleration curve shows that the magnitude of the acceleration caused by the velocity oscillation was about 0.5 m/s (0.051 g). The small oscillation in velocity was caused by the velocity controller's inability to follow the velocity commands with exact precision. The time domain plot also shows that the velocity dipped slightly at 1423

seconds. This was due to a downshift by the automatic transmission (we do not control transmission gear changes in our present experiments). The algorithm compensated by gently accelerating the host to a slightly higher speed so that the headway continued to decay to the desired headway. The host velocity then decreased to match the target vehicle's velocity once the desired headway was reached.

The deceleration profile shows that the maximum deceleration occurred at slightly below –1.0 m/s. It is apparent that light braking was applied during this deceleration (we did not track when brakes were applied since this is the function of the velocity controller). The deceleration then dropped to 0.5 when throttle was used to slow the host vehicle. The lack of large jumps in the acceleration curve supports the qualitative observation that the deceleration was "smooth". The maximum virtual force applied to the host was –1.7 N (the force was normalized by diving by the mass of the truck so that it more conveniently fit on the plot). It is apparent when comparing the force and acceleration curves that there is a delay between the application of a force and the resultant acceleration of the host. This is caused by the velocity controller's time constant.

The range vs. range rate and time domain data for a higher speed ACC experiment is shown in Figure 3.5 through Figure 3.7. First, notice that the range sensor was fairly noisy in this experiment. This may be due to the fact that the road surface of the main line was in worse condition than the low volume road, where the low speed experiments were performed. The noise in the range signal is amplified when it is differentiated to obtain the range rate (for the laser sensor). This accounts for the lone points to the right of the switching line. Regardless, the algorithm proved to be rather robust to noisy range data as shown here and by simulation in the previous chapter. The desired headway time was increased to 2.0 seconds for these experiments for no reason other than to provide the truck driver more headway distance. The host set speed was 25 m/s (90 kph) and the target was traveling at 18 m/s (64.8). The range vs. range rate plot shows that the target was in the linear force region for the duration of the experiment (excluding the range sensor noise). The trajectory of the target state (R, $\dot{R}$) in this phase plane was mostly linear as desired. The time domain plot shows that the host velocity and range decreased

smoothly to their desired values. The acceleration curve shows that the deceleration was small, which supports the qualitative observation that the deceleration was mild. The range sensor noise shows up in the force curve. Its effect on the truck's acceleration is damped by the longitudinal admittance controller and the impedance of the truck.

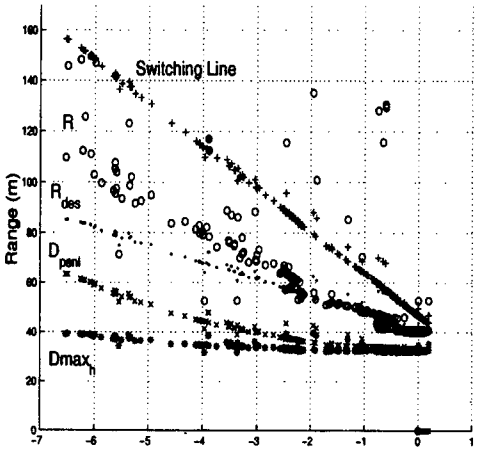

Figure 3.5: R vs. $\dot{R}$ plot, high speed ACC functionality with Vtar = 18 m/s

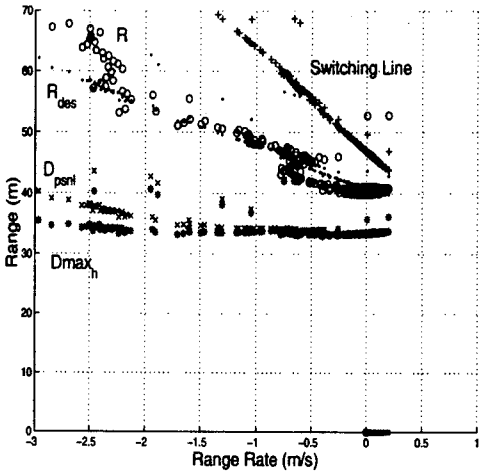

Figure 3.6: Rvs. $\dot{R}$ plot, high speed ACC, scaled to show the end of the maneuver

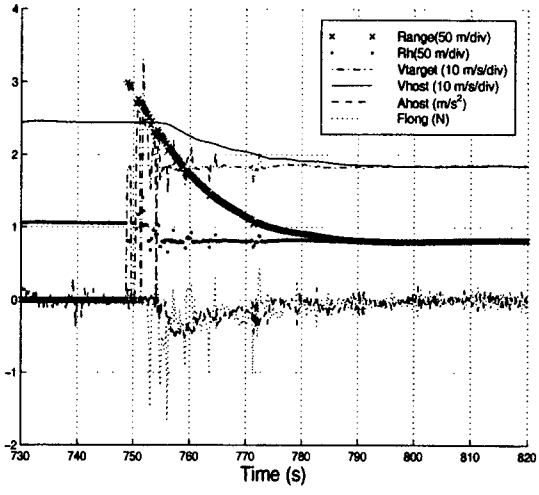

Figure 3.7: Time domain plot, high speed ACC functionality with Vtar = 18 m/s

It was observed through various experiments that the velocity oscillation experienced in the low speed experiments disappeared at higher speed. This is due to the diminished effect of DGPS noise on the velocity and acceleration calculations at higher speed and the velocity controller more precisely following the command velocity produced by the virtual bumper longitudinal controller. Furthermore, transmission gear changes were less frequent at higher speeds and produced fewer velocity shifts. In fact, it was observed that it is more important to the performance of the longitudinal controller for the velocity controller to accurately follow velocity commands than it is for the sensor data to be free of noise.

The ACC functionality of the headway controller proved to work well at both lower speed and higher speed scenarios. The system has some small velocity oscillations at lower speed (below 2 m/s (7.2 kph) to 1 m/s (3.6 kph) due to the velocity controller's inability to follow the velocity commands with high fidelity at such low speeds. Vehicle speed controllers are not usually designed for this task. At higher speeds, however, the performance of the velocity controller improves and the oscillation is eliminated. It was also shown that the algorithm is fairly robust to noisy range data.

3.5 Major Traffic Slow Down Experiments

In this scenario, the goal is to test the non-linear forces and the transition from the non-linear to the linear force region. The target is in the non-linear force region when it is detected at the maximum sensor range and the target state $(R, \dot{R})$ is below the non-linear personal space boundary.

We produce this traffic scenario by having the host travel at a velocity significantly greater than the target vehicle so that the host vehicle must apply brakes to reach the desired headway and velocity. We also want to experimentally verify if the transition from the non-linear to the linear force region has a smooth deceleration curve.

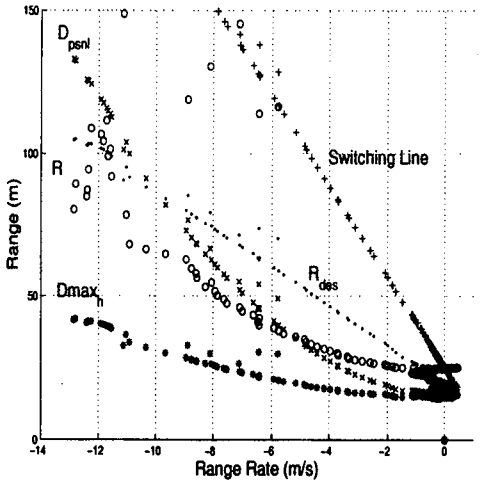

Figure 3.8: R vs. $\dot{R}$ plot, major traffic slow down scenario with Vtar = 8 m/s

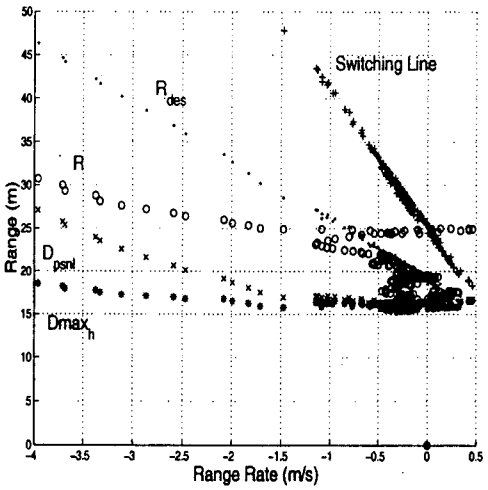

Figure 3.9: R vs. $\dot{R}$ plot, major traffic slow down, scaled to show the end of the maneuver

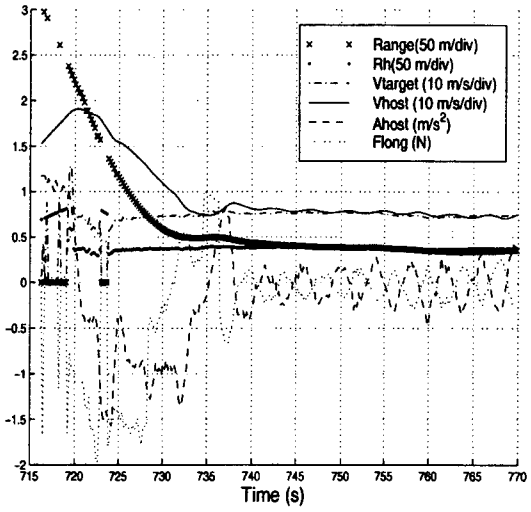

Figure 3.10: Time domain plot, major traffic slow down scenario with Vtar = 8 m/s

We first present the result of an experiment in which the target was moving at 8 m/s (28.8 kph) and the host was initially at 18 m/s (64.8) before control action was taken. Figure 3.8 through Figure 3.10 show the range vs. range rate phase diagrams and the time domain plot respectively. The target was first detected at long range in the linear force region. These data points (range

sensor noise) are few and occur at long range. They do not significantly affect the host's velocity as it attempts to reach the set cruising velocity. The target then is detected in the non-linear force region where the forces become large. The host undergoes a fairly constant deceleration as indicated by the parabolic trajectory in Figure 3.8. When the $(R, \dot{R})$ trajectory reaches the linear force region, it is below the desired trajectory (R_{des}) line. Therefore, the forces remain negative although they are reduced in magnitude. The forces continue to decrease as the $(R, \dot{R})$ data points approach the desired trajectory. The forces become positive when the $(R, \dot{R})$ data points overshoot the desired trajectory line. The delay in the truck responding to the velocity controller's adjustment in throttle/brakes caused the host to temporarily overshoot the desired headway. As the host's velocity dips below the target velocity, $\dot{R}$ becomes positive and the target leaves the personal space. Cruise control forces are applied to the host to increase its velocity and decrease the headway so that the target re-enters the personal space. Then, the linear forces bring the target to the desired headway.

The transition from the non-linear force region to the linear force region is acceptable when the acceleration profile is considered. The deceleration decreased from around 1.0 m/s^2 to 0.5 m/s^2 in a smooth manner as the target made the transition from non-linear to linear forces. The velocity undershoot was caused by the time constant of the velocity controller. The forces became positive as the $(R, \dot{R})$ trajectory passed the desired trajectory, but there was insufficient time for the velocity controller to adjust the speed of the host before $\dot{R}$ became positive. Although this behavior was not expected, it errors on the side of caution. This is a relatively severe driving maneuver and the fact that host velocity decreased a little too rapidly is acceptable. Since this is a headway controller, the important fact is that the host did not significantly undershoot the desired headway.

The result of an experiment with the host initially traveling at 25 m/s (90 kph) and the target at 15 m/s (54 kph) is shown in Figure 3.11 and Figure 3.13. From the range vs. range rate phase plot it is clear that the target was first sensed consistently in the non-linear force region. There is some scattered target measurements of which only four are under the switching line. These

points are noise as can be seen by their corresponding switching line, D_{psnl}, R_{des}, and $Dmax_h$ points which do not correspond with the "clean" curves generated by the less noisy $(R, \dot{R})$ points. One can see from the figure that the host performed a constant deceleration (parabolic trajectory) until it reached the non-linear force region. Then, the linear forces were applied which brought the host to the desired headway.

An analysis of the time domain plot reveals that the host velocity and headway both decrease to their desired values with little or no overshoot. The acceleration curve shows the larger deceleration caused by the non-linear forces followed by a smaller deceleration caused by the linear forces until the accelerations becomes practically zero (plus noise). It is also shown that the forces drop significantly after the transition from non-linear to linear forces and actually become positive for a brief time. The positive force slows the deceleration until the target state $(R, \dot{R})$ is near the desired range-range rate trajectory. The forces then become small as the target approaches the desired state.

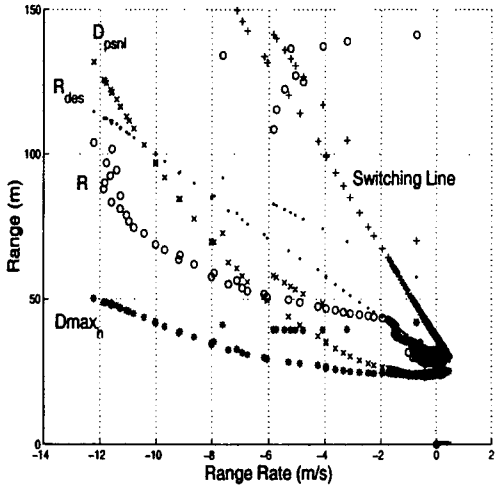

Figure 3.11: R vs. $\dot{R}$ plot, major traffic slow down scenario with Vtar = 15 m/s

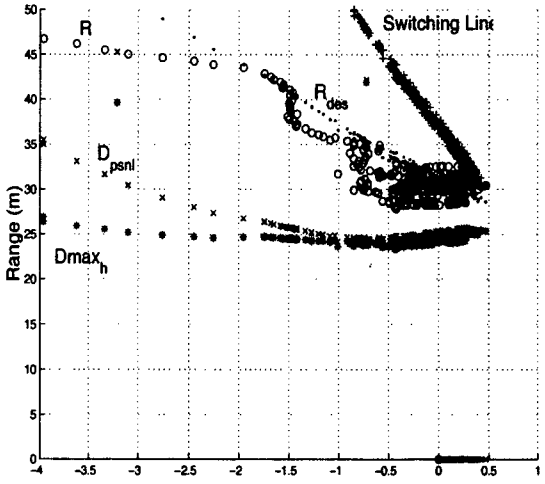

Figure 3.12: R vs. $\dot{R}$ plot, major traffic slow down, scaled to show the end of the maneuver

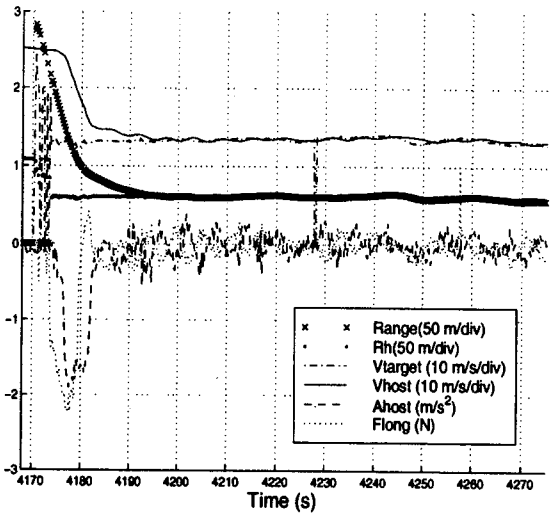

Figure 3.13: Time domain plot, major traffic slow down scenario with Vtar = 15 m/s

It is interesting that in the first experiment shown for this driving scenario, the host velocity undershot the target velocity before settling to the desired headway. In the second experiment the host velocity remained above the target's velocity so that $\dot{R}$ stayed negative down to the desired headway. This was an unexpected result based on the computer simulations. Whether the host will undershoot the target's velocity seems to depend upon the state at which the $(R, \dot{R})$ trajectory enters the linear region. InFigure 3.8, the target enters the linear force region with a range of around 37 m and a range rate of –5.2 m/s. Upon the transition to the linear force region, the forces drop to around –0.5 N (normalized) and decrease for four seconds and become positive. In Figure 3.11, the target enters the linear force region at a range of around 52 m and a range rate of –6.1 m/s. When the forces become linear, they drop to –0.4 N very briefly and then become positive. This quicker transition to positive force allows the velocity controller enough time to decrease the deceleration of the host and to achieve the desired trajectory. Although the initial conditions for both experiments in the range vs. range rate phase diagram are very similar, the range data in Figure 3.11 is less noisy which helped decelerate the host vehicle sooner to force the $(R, \dot{R})$ trajectory into the linear force region at a higher range. Another difference in the experiments is the final steady state host/target velocity. The target velocity in the first

experiment was lower. It was generally observed that the velocity controller more accurately followed velocity commands at higher speed. Perhaps this contributed to the velocity undershoot of the first experiment.

The transition from the non-linear to the linear force region is similar to the simulations in the sense that the acceleration change is mild. We did observe that the host velocity became less than the target velocity at a range greater than the desired headway for some experiments. This seems to be related to the initial conditions of the host state (R, $\dot{R}$) as it enters the linear force region in the range vs. range rate phase plot. This state is affected by the initial conditions of the target when it is first sensed as well as sensor noise and drop outs. The final result in both cases is that the host approaches the target at the desired headway with little or no overshoot.

3.6 Critical Stop Experiments

To test the collision avoidance capability of the virtual bumper longitudinal algorithm in an emergency situation, we conducted an experiment in which the target was stationary. This simulates traffic scenarios such as severe traffic jams, approaching a stalled vehicle or a stoplight.

In the first critical stop experiment, the host was initially traveling at 10 m/s towards the stationary target. This scenario is similar to a vehicle approaching a traffic light at a city street intersection. This slow speed was chosen to test the linear forces and their ability to stop the target. The result is shown in Figure 3.14 through Figure 3.16. The target is first detected in the linear force region and the linear forces slow the host down and the (R, $\dot{R}$) trajectory converges to the desired state. At around –2 m/s, the host maintains a constant velocity until the range decreases to the desired headway at which point the stopping force stopped the vehicle. Note that this is not exactly the trajectory the controller was designed to pursue.

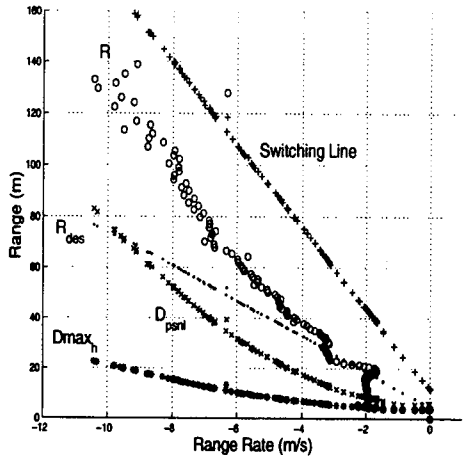

Figure 3.14: R vs. $\dot{R}$ plot, critical stop scenario with Vtar = 0 m/s, Vhost initially at 10 m/s

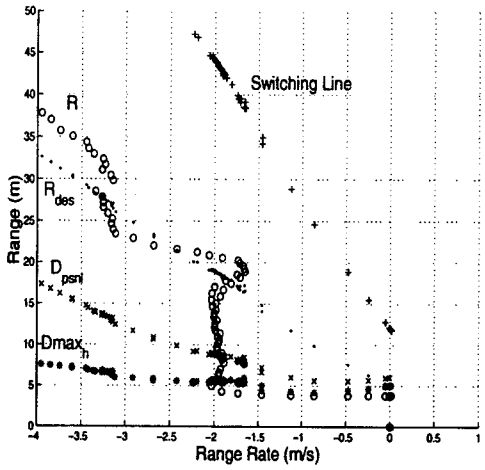

Figure 3.15: R vs. $\dot{R}$ plot, critical stop scenario, scaled to show the end of the maneuver

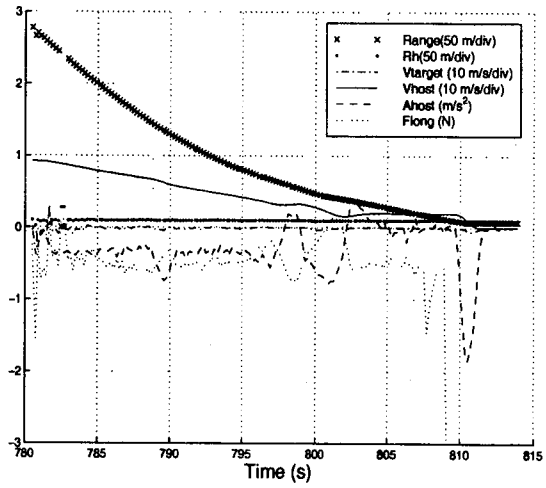

Figure 3.16: Time domain plot, critical stop scenario with Vtar = 0 m/s, Vhost initially at 10 m/s

The time domain data shows that the acceleration is fairly constant while the linear forces are applied. Then, the acceleration drops to zero as the host coasts toward the target at 2 m/s.

Finally, the acceleration peaks when the brakes stop the target vehicle. The final range is less than the desired headway range. This is due to the characteristics of the braking system.

The braking system consists of a linear actuator that pulls a cable attached to the brake pedal. Although this system is able to apply full brake actuation to stop the Navistar tractor, there is insufficient sensitivity to allow very small velocities to be realized. This limitation made it difficult for the virtual bumper longitudinal controller to follow desired velocities at low speeds. The virtual bumper controller produces velocity commands in all ranges down to zero and assumes that the velocity controller is able to follow the commanded velocity. This assumption is not quite the case when operating close to zero m/s. We were, however, able to work around this hardware limitation by modifying the algorithm slightly. The idea is to determine the lowest velocity command the velocity controller is able to maintain (V_{min_h}). Then, generate a zero velocity (stop command) only when it is time to stop the host. This was accomplished by letting the virtual forces slow the host down to the minimum allowed velocity (2 m/s) and then apply a stopping force ($F_{stop} = -Dmax_h$) when the range decreases past the minimum headway (R_{Ho}) for a non-moving target.

The stopping force causes the velocity controller to apply the brakes at full stop, which does cause a sudden stop. This is not entirely comfortable for the passengers, but it is the only way to ensure a full stop given the current hardware. Prior to the implementation of this stopping force, the linear forces would drive the velocity command below V_{min_h}. This caused the host to stop before the desired headway because the brake actuator could not achieve the correct brake position needed to slow the host to the commanded velocity. The virtual forces then became positive in order to attempt to achieve the desired headway. This caused the host to move forward at V_{min_h} (brakes off, throttle idle), which is faster than the command velocity. Negative forces would ensue and the cycle would repeat itself until the range became less than R_{Ho}. The host, in effect, would "stop and go" until the desired headway was reached. This certainly is a less desirable response than the one shown in Figure 3.16.

The next critical stop experiment was conducted with the host initially traveling at 18 m/s (39 mph). The result in Figure 3.17 through Figure 3.19 shows that at first, the radar detected a false target at 180 meters. This false reading caused virtual forces to be applied to the host, which slowed it down slightly. The true target is then sensed in the non-linear force region (range ~ 180 m, range rate ~ -17 m/s) and non-linear forces are applied that cause a somewhat constant deceleration of the host. Notice that the deceleration does increase toward the end of the stopping maneuver. This is caused by the brake actuation system, which cannot achieve the low velocity as explained previously.

This experiment highlights an important factor that affected the longitudinal controller. False targets proved difficult for this system because there are no preventative routines to deal with it. While the algorithm is robust towards sensor noise, it assumes that the target data is legitimate in the sense that it represents a real target vehicle. No attempt to verify this information is provided.

No experiments with larger host velocities were attempted because our range sensors were limited in range. The Optech range finder had a longer maximum range, but the range reading at longer distances proved inaccurate. Differentiation to obtain range rate verified this observation. The range rate signal was so noisy at longer ranges that the maximum range of the Optech had to be reduced by filtering. This gave the Eaton VORAD radar the longer maximum range; however, the one-dimensional nature of this sensor caused it to detect false targets at longer range. The limiting of the maximum useful range of our sensors prohibited us from conducting critical stop maneuvers at higher speeds. The algorithm and velocity controller were capable of handling this task if knowledge of the target is known at longer distances.

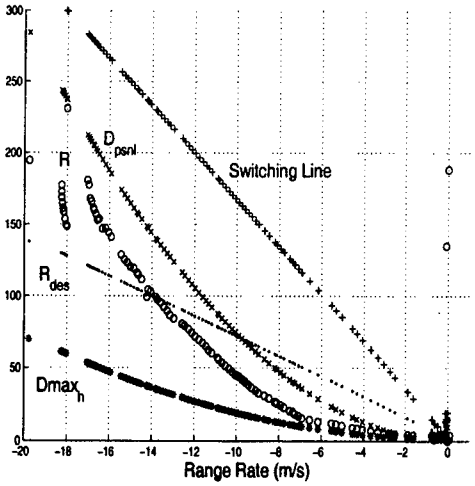

Figure 3.17: R vs. $\dot{R}$ plot, critical stop scenario with Vtar = 0 m/s, Vhost initially at 18 m/s

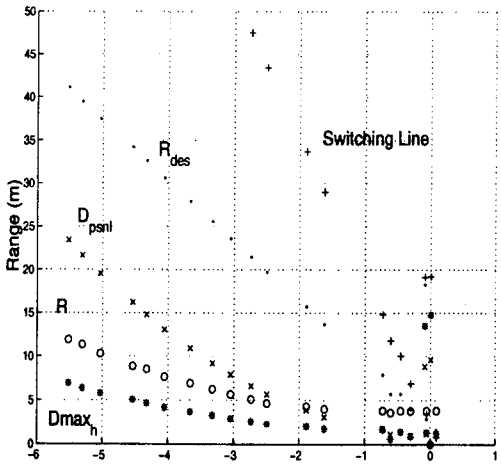

Figure 3.18: R vs. $\dot{R}$ plot, critical stop scenario, scaled to show the end of the maneuver

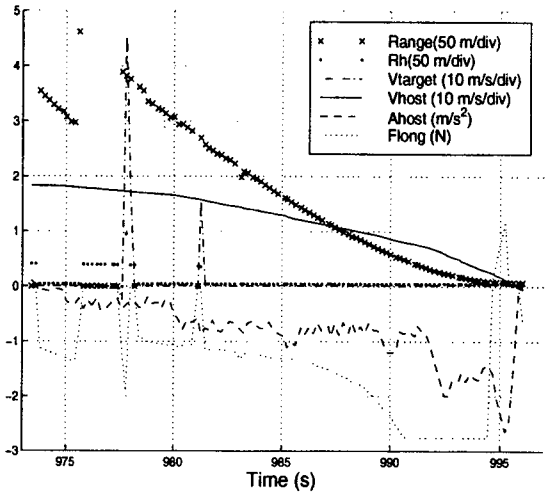

Figure 3.19: Time domain plot, critical stop scenario with Vtar = 0 m/s, Vhost initially at 18 m/s

3.7 Miscellaneous Experiments

We have shown that the algorithm works well in maintaining a constant headway time for a target vehicle moving at a constant velocity. It is equally important that it can track an

accelerating/decelerating target in a similar manner as a human driver would. The experiments in this section test this ability.

It should be noted that in this section, we do not show the range vs. range rate phase plot. These plots are more difficult to understand than the ones in previous experiments due to the varying trajectories caused by each scenario. Detailed explanations of each $(R, \dot{R})$ trajectory would be required in order to understand what happened for each experiment. The time domain plots adequately display what took place in the experiment and was found to be easier to read.

3.7.1 Target Cuts into the Host's Lane of Travel

The 'cut-in' scenario was tested by having the target cut into the lane of the host and then decelerate. This situation is very common in highway driving and we wanted to verify that the host would not produce large decelerations that would cause the system to 'overreact' to this frequent driving situation. Figure 3.20 shows the time domain plot of this experiment. The target is detected with a positive relative velocity and a range well below the desired headway. Thus, negative virtual forces are applied to the host, which cause it to decelerate. The low deceleration confirms that the system did not overreact to the target, but simply decreased the throttle, much the way a human driver would. The headway increases towards the desired range, but does not achieve the desired headway only because the experiment was terminated due to host approaching the end of the test track.

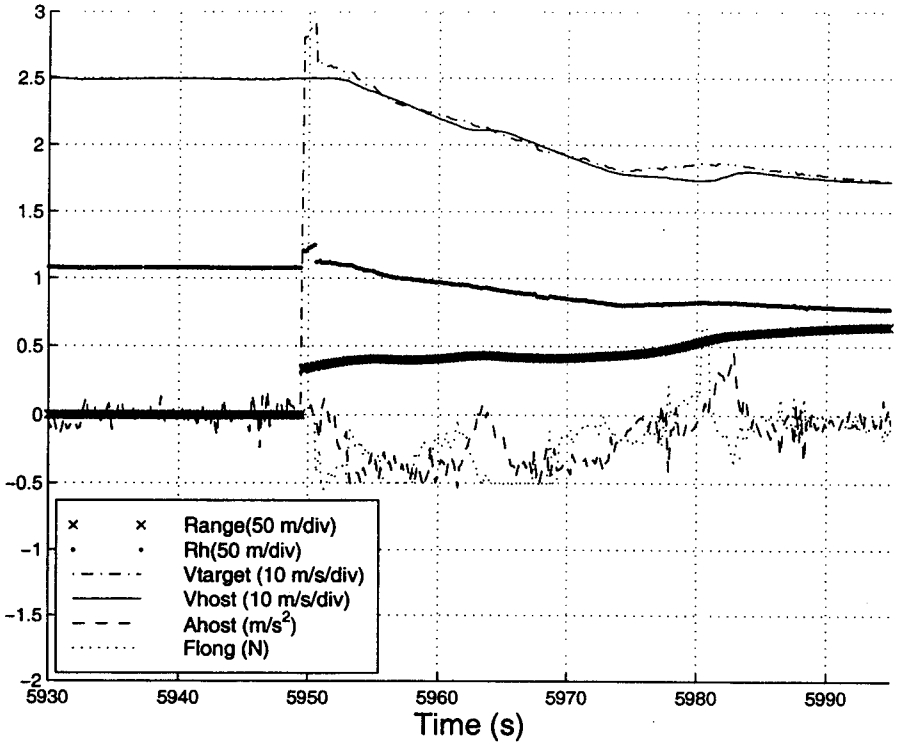

Figure 3.20: Time domain plot, cut-in scenario with Vtar initially at 25 m/s

3.7.2 Target Accelerates After Being Tracked by Headway Controller

To demonstrate the effect of the cruise control force, we performed an experiment in which the host matches the speed of the target, then the target accelerates. We want to demonstrate that the host can accelerate to follow the target as long as the target is traveling below the host's desired traveling speed (as set by the driver).

The cruise control force was developed so that the host always travels at the set cruise control speed when a target is not in the personal space boundary. To test this, we let the host decelerate to match the velocity of the slower moving host and then accelerated the target vehicle. The result is shown in Figure 3.21. The range sensor detects the slower moving target and decelerates the host to match the target's velocity and achieve the desired headway. At about

3372 seconds, the target accelerates. This increases the headway since the target vehicle is now traveling at a higher speed than the host. The cruise control forces are applied (positive) to accelerate the host to match the target's velocity. The desired headway also increases as the host accelerates, but the actual headway increases more because the target is now traveling at a higher speed than the host. As the host velocity increases, the error in headway decreases, but does not get to zero. This is simply because the host approached the end of the test track and the experiment had to be terminated.

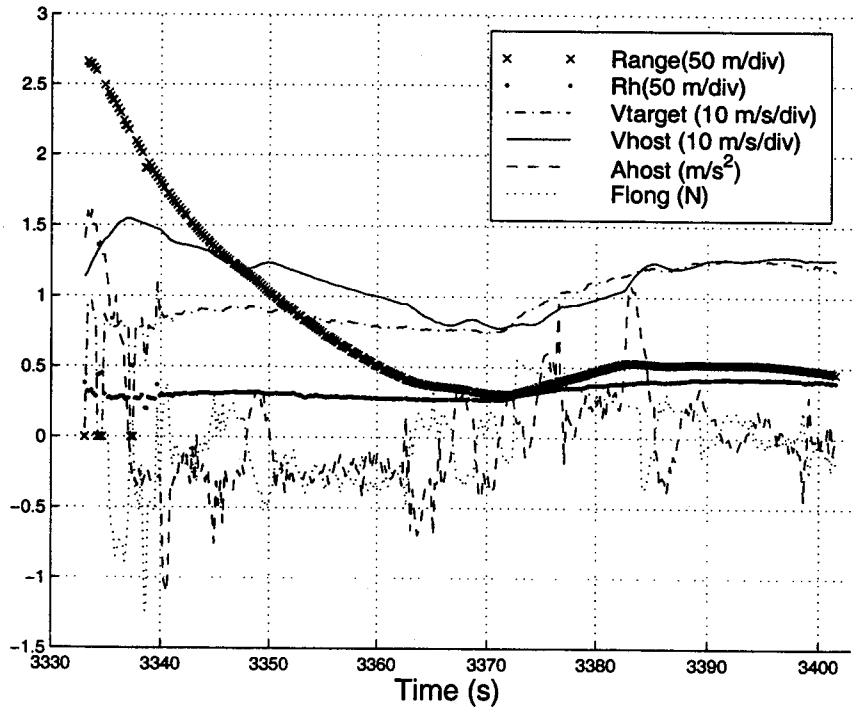

Figure 3.21: Time domain plot, target accelerates after host has reached desired headway and velocity

It is clear that the cruise control forces effectively accelerate the host when the target moves out of the personal space boundary and the host is traveling below the desired speed. The cruise control forces ensure that the host will increase its speed to the desired cruising speed if the target moves out of the personal space boundary by accelerating or changing lanes.

3.7.3 Target Decelerates After Being Tracked by the Headway Controller

This experiment was designed to determine if the longitudinal controller can react quickly enough to a target that has been followed under ACC but suddenly decelerates. This situation is fairly common when there is a traffic slowdown in tight traffic.

This scenario was tested by having the host approach a target vehicle traveling at 10 m/s. After the host slowed down to match the target vehicle velocity, the target then came to a stop. This is shown in Figure 3.22. The desired headway and actual headway decrease as the target vehicle slows down. The virtual damper reacts to the velocity difference and applies a virtual force to slow the host. The final headway ended up below the desired headway. This again is due to the fact that the velocity controller could not follow the small velocity commands. Only when the range is less than the minimum headway range at zero velocity (R_{Ho}) is the stopping force applied.

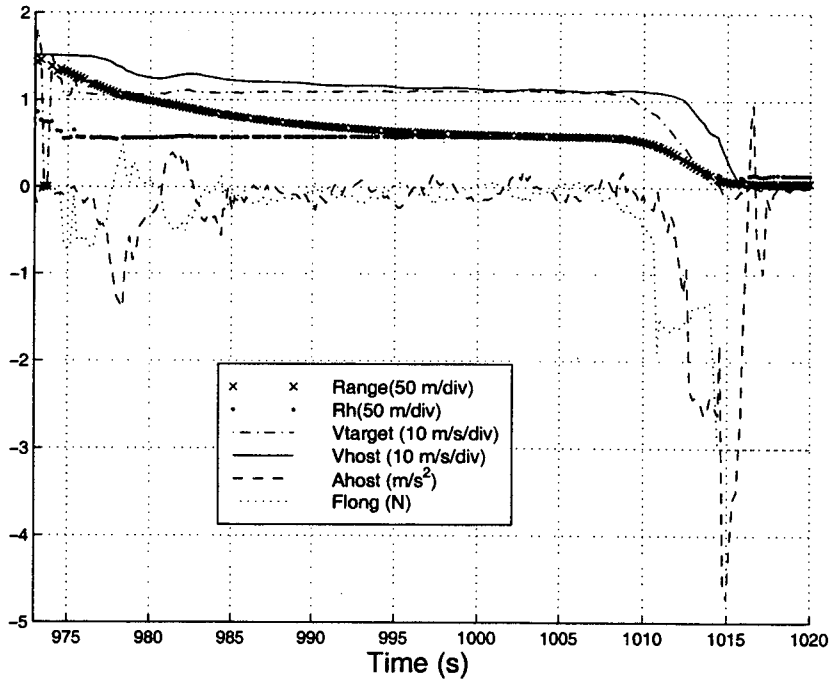

Figure 3.22: Time domain plot, target coming to a stop after host established a steady state headway

3.7.4 Stop and Go Traffic

The virtual bumper is a collision avoidance algorithm that is designed to mitigate or avoid collisions. While stop and go traffic is not a critical application of a collision avoidance system, it was in the interest of thoroughness to test the algorithm in this scenario.

To simulate stop and go traffic, the target was driven at speeds between 0 and 5 m/s. The result is shown in Figure 3.23. The host successfully followed the target. Notice that the desired headway was not tracked precisely. This is because the target leaves the personal space as its speed increases. The cruise control forces then accelerate the host until the target is once again within the personal space boundary. Then, the headway is reduced to the desired value. However, the target decelerates before the desired headway is achieved.

The spikes in the acceleration curve are caused by the transmission shifts and the brake hardware limitations. The transmission shifts cause a jump in positive acceleration at low speeds. This is compounded by the fact that we did not have a trailer hooked up to the tractor. A trailer with large mass acts to damp the acceleration due to transmission gearshifts. The brake hardware limitation was documented in the previous section.

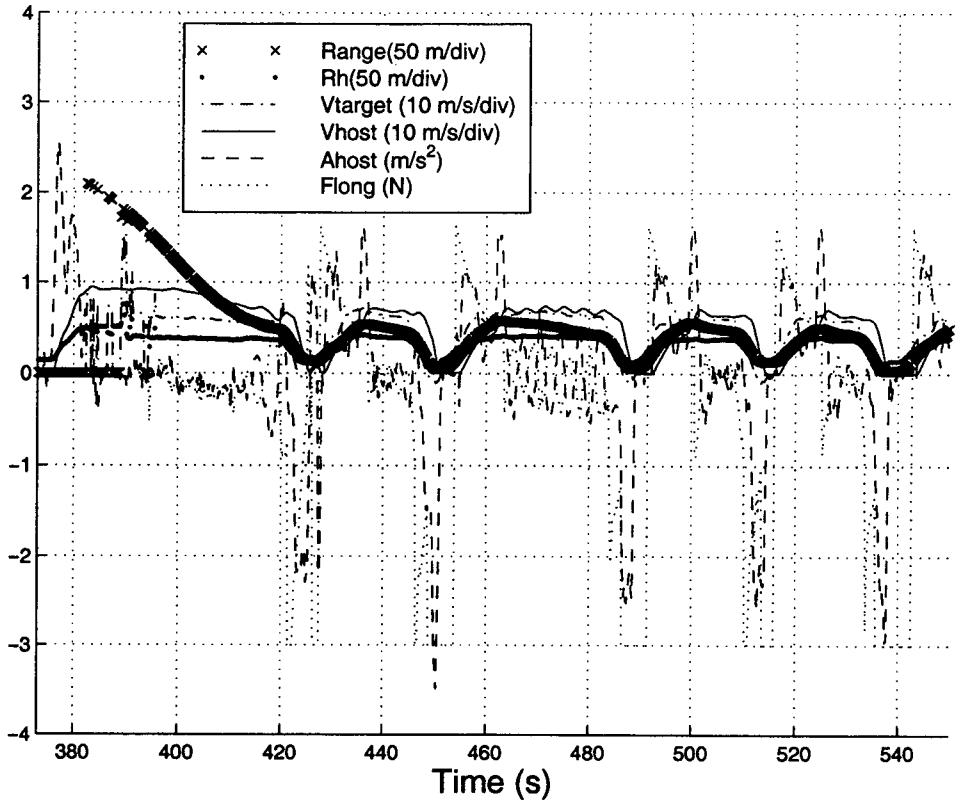

Figure 3.23: Time domain plot, stop and go traffic scenario

We have demonstrated that the virtual bumper longitudinal algorithm is able to control the headway distance to the proceeding vehicle in both normal and emergency driving situations. The range sensors and brake/throttle hardware caused some behavior that was not predicted in the simulation, but minor changes to the algorithm overcame these difficulties. The next chapter details the conclusions of the experimental results and makes recommendations on how to improve the system in the future.

Chapter 4

Conclusions and Recommendations

4. Conclusions and Recommendations

We have thus far developed the theory of the virtual bumper longitudinal controller, illustrated results in simulation and showed the experimental results of the system implemented on a Navistar tractor. In this chapter, we summarize the findings of this report as well as discuss what our recommendations are for future work to improve the current system and to test the other parts of the virtual bumper algorithm.

4.1 Simulation and Experimental Results

In this section, we describe some of the differences between the simulated and actual longitudinal controller. We will explain our hypothesis as to why the two are different and detail the strengths and weaknesses of the algorithms for each driving situation.

4.1.1 Conclusions of ACC Functionality

In the simulation, the range sensor data was perfect and the range sensors had a maximum range well above the personal space boundary. Thus, the algorithm slowed down the host to exactly the speed of the target and the desired headway. The system implemented on the Navistar tractor achieved the desired headway, but displayed some oscillation at lower speed due to the tracking error of the velocity controller caused by the braking actuation methodology.

For the ACC experiment, the host approached a slightly slower moving vehicle to test the ability of the linear forces to slow the host to the desired velocity and headway. This was done for two different target speeds; 9 m/s and 18 m/s. The headway controller successfully slowed the host to match the target's velocity and the headway decreased to the desired value with little or no overshoot.

It was observed that the steady state velocity/headway had a slight oscillatory behavior during the slower speed experiment. This was attributed to DGPS error and the velocity controller, which could not control the speed as accurately at lower speed due to transmission gear shifts and a general decreased throttle sensitivity. At higher speed, the oscillations in velocity disappeared.

It was also noticed that the personal space boundary needed modification so that the target did not "drift" in and out of the personal space boundary. This occurred because the personal space boundary converges with the desired headway when $\dot{R}$ goes to zero. Any slight positive deviation in relative velocity would cause the target to leave the personal space boundary. In simulation this was not an issue because the target velocity stayed constant, the host velocity controller maintained the exact commanded velocity and the range sensors perfectly measure the range and range rate to the target vehicle. As expected, in the real world, all these assumptions are not accurate. Velocity deviation in the target vehicle, tracking error in the velocity controller and range sensor error caused the target to slightly deviate from the steady state headway.

When the target leaves the personal space boundary, virtual cruise control forces are applied which push the host forward. It was discovered that these forces, larger than the linear forces near steady state, would decrease the headway below the desired value. This brought about a cycle of oscillation in headway. The addition of a headway buffer to the personal space prevented this behavior because a slight deviation from headway is adjusted for by the linear forces, which are smaller near the desired target state (R, $\dot{R}$). The linear forces gently adjust the host's velocity to bring the headway back to the desired state.

The deceleration during the ACC experiments was generally gentle. The linear forces were designed to provide deceleration due to throttle adjustment or light braking. Thus it was no surprise that the deceleration was small.

It was also observed that the algorithm is rather insensitive to range sensor noise, at least at long range. The Optech range finder was noisy at long ranges, but this didn't seem to affect the system. More detail on hardware problems will be given in §4.2.

4.1.2 Conclusions of Major Traffic Slowdown Scenario

Schiller [2] demonstrated the controller's ability to regulate the headway when the host is initially traveling significantly faster than the target. In this type of traffic scenario, moderate to

heavy braking is needed to slow the host. This tests the non-linear forces and the transition from the non-linear force region to the linear force region. Schiller's simulation showed that the longitudinal controller was able to successfully handle this scenario and that the transition between force regions was smooth.

Analysis of our experiments demonstrated the same result. The host successfully decreased its velocity to match the target and brought the headway to the desired value. The deceleration curve showed that the transition between force regions did not exhibit large deceleration jumps. As expected, the deceleration did drop when the forces became linear.

This experiment was performed at two different velocities. The higher speed experiment was very much like Schiller's simulation. The lower speed experiment, however, displayed an unexpected behavior. The host velocity undershot the target velocity during the transition between non-linear and linear forces.

It appears that the time constant of the velocity controller caused this behavior. The forces became positive as the target state (R, $\dot{R}$) approached the desired trajectory line (R_{des}). However, the delay between the velocity controller receiving the velocity command and changing the host velocity caused the host to slow too much. The target left the personal space boundary and the host sped up and completed the headway maneuver.

We believe this happened in the slow speed experiment and not in the higher speed experiment due to the initial conditions (R, $\dot{R}$) when the target enters the linear personal space. In the higher speed experiment, the target state was at a greater range and lower range rate. As the linear forces became positive, the velocity controller had more time to change the host velocity because the velocity command started increasing at a lesser $\dot{R}$ (more to the left of the range vs. range rate phase plot). It seems that how the target enters the linear personal space determines whether this velocity undershoot occurs.

Range sensing plays a critical role in how the target state enters the linear force region. The longer the maximum sensing range, the sooner brakes are applied and the sooner the target state enters the linear force region. This increases the time that the velocity controller has to reduce the deceleration. The resultant trajectory on the range vs. range rate phase plot is smooth with no velocity undershoot.

Since this is a rather severe maneuver and since the unit was designed as a headway controller and not explicitly a velocity controller, the experimental results were acceptable. It is better to slow down too soon and before the desired headway than to decelerate too late and risk a collision.

4.1.3 Conclusions of the Critical Stop Scenario

We define the critical stop scenario as one in which the target vehicle is at rest. The host approaches the target at different speeds to determine the longitudinal controller's ability to stop the host. Schiller did not simulate this type of driving situation with braking alone. He did simulate this driving scenario with the full virtual bumper so that the host changed lanes instead of coming to a stop. If a vehicle is in the adjacent lane, a lane change may not be possible and only braking would then prevent/mitigate a collision.

This scenario was most sensitive to the shortcomings of the braking hardware. More detail will be provided in §4.2, but to summarize, the velocity controller could not achieve small velocity commands due to the brake hardware's inability to adjust the brakes with enough sensitivity. This resulted in a dead band in the possible velocity commands that could be given to the velocity controller. A command of zero would stop the host and a command greater than zero and less than V_{min_h} would produce a velocity of V_{min_h}. This minimum velocity turned out to be the coasting speed (throttle was idling, brakes not applied) of the host, 2 m/s.

A slight modification to the algorithm was necessary to work around this hardware limitation. A stopping force (F_{stop}) was applied when the range became less than the minimum desired headway, R_{Ho}. When the velocity command was below V_{min_h}, the host would travel at V_{max_h}

until the stopping force was applied. This caused an abrupt stop at the end of the maneuver, which is not entirely comfortable for the driver. Given the hardware limitations, this was an acceptable solution.

The low speed experiment showed that the linear forces were sufficient to slow the host down to V_{max_h}. Then the braking forces stopped the host at a range slightly less than the minimum headway, R_{Ho}. The deceleration during this experiment was low and fairly constant until the brakes were finally applied to stop the tractor.

The higher speed experiment involved the host initially traveling at 18 m/s. The host decelerated with a constant value until the end of the maneuver when the full brakes were applied to stop the host. The controller performed very well under this scenario as the non-linear forces effectively produced a constant deceleration.

Experiments at higher speeds were not achievable because the maximum range of the range sensors was limited. At higher speeds, knowledge of the target vehicle's range and velocity are needed for the force to be applied early enough to stop the host. We were also limited by the velocity controller's maximum deceleration, D_{max_h}, which was experimentally determined to be 2.76 m/s.

4.1.4 Conclusions of Miscellaneous Driving Scenarios

This category of driving scenario involves a target that is not traveling at a constant velocity. Real highway traffic is dynamic, the mean flow of traffic changes continuously due to the number of lanes, congestion, construction, etc. We tested the controller's ability to maintain the desired headway when the target decelerates, accelerates, or both.

The first driving scenario tested is called the 'cut-in'. This happens when a target switches to the lane in which the host is traveling and is within the desired headway of the host. The target then slows down. This experiment was performed and the results were favorable. The host gently

decelerated along with the target and the headway, which was originally below the desired headway upon target acquisition, increased to the desired value.

The next driving scenario involved the target accelerating while being followed by the host under ACC. This was designed to test the cruise control forces and their ability to accelerate the host. The results were again as expected. As the target accelerated, the cruise control forces were applied and the host accelerated to match the speed of the target.

The target vehicle deceleration experiment also showed acceptable results. The target was followed under ACC and then decelerated to a stop. The virtual forces decelerated the host to a stop, but the final headway was less than the desired minimum headway. This happened because of the brake hardware limitation that was previously documented.

Even though the virtual bumper was designed to be a collision avoidance algorithm, it was interesting to observe whether it could operate in stop and go traffic. This is not a critical driving situation as the speeds are low, but we wanted to determine if the system is stable for a target that accelerates and decelerates repeatedly.

The virtual bumper longitudinal controller was successful in maintaining a safe headway in this traffic scenario. The headway did get significantly larger than the desired headway as the target accelerated. This was due to the cruise control forces, which were limited in magnitude to provide moderate accelerations. Our tractor did not have a trailer, but a loaded trailer prohibits large acceleration. This behavior may be unavoidable, but human drivers do not pursue headway distances with great accuracy. Fancher studied the headway times of human drivers and cars under ACC. He found that the standard deviation of headway time was greater for human driving than for ACC [36]. A tradeoff must be made between strict following of headway distance and driver comfort. We chose as part of our design goals to emphasize driver comfort when necessary. Thus, positive deviations in headway are acceptable for this tradeoff.

4.2 Hardware Limitations

The virtual bumper depends upon the hardware to sense the surrounding vehicle and to maneuver the host as commanded. Acceleration limits were considered so that the output commands could be achieved by the lower level vehicle controllers. We found that this was not possible with the present hardware aboard the Navistar tractor test bed.

The brake system consists of a linear actuator connected to a cable and pulley system that pulls the brake pedal down. Stiction in the actuator and a small length of travel prohibited precise control of the brakes. The result was that the velocity controller could not produce small changes in velocity. The brake system was adequate for emergency stopping.

We also discovered that the brake system behaved differently on different days. Temperature seemed to be a key component in the various behavior. Cold days were especially problematic, as we had to increase the voltage to the brake actuator to even move it. This forced us to have to tune the braking system before we could perform experiments. If the brakes were out of tune, they would tend to 'grab' when even light braking was desired. This caused the host to slow down more than the commanded velocity. The headway controller would compensate by increasing the velocity of the host. When brakes were eventually required to slow the host, they would grab again and the cycle would repeat itself down to the desired headway.

The throttle was controlled by a PWM signal to the engine. At high speeds, the throttle controller worked very well. At lower speeds, the throttle controller was affected by gearshifts. The gear shifts introduced a jump in velocity (disturbance) which had a tendency to induce velocity oscillations. Despite this, the virtual bumper longitudinal controller was robust enough to the disturbances caused by the gear shifts and adjusted the command velocity in a manner that reduced velocity oscillations to a low amplitude.

In general, it was observed that it is critical that the velocity controller follow the velocity commands of the virtual bumper. We considered the delay of the velocity controller by adding

the predictive term T, but we assumed that the velocity commands would be followed with the assumed phase delay. It was discovered that tracking errors (besides the expected phase delay) caused oscillations in velocity.

The range sensors posed several problems. The Optech laser range finder was noisy at long range. Filters were designed to improve the data from the sensor, but they also reduced the maximum sensing range. The Eaton VORAD radar worked well for the critical stop experiments, but was not useful in the ACC experiments. Since this is a Doppler sensor, the data became unpredictable as the relative velocity between vehicle decreased. Also problematic were false targets. The radar does not provide an azimuth angle to the target so false targets at long range are common. The four degree cone angle causes the radar signal to spread to adjacent lanes and past the road boundary. This proved very problematic for the longitudinal controller. While robust to sensor noise, the algorithm depends upon the range sensors measuring real targets within the personal space boundary. Since we do not know where the targets are located laterally, we had to limit the maximum range of the sensor to filter out false targets.

4.3 Recommendations

The virtual bumper longitudinal controller has proven to be effective in controlling the headway to the preceding vehicle. In numerous driving situations, the controller was able to bring the headway to the desired value as long as the sensor data was good and the velocity controller was able to achieve the command velocity commands. We believe that the most significant improvements can be made by improving the sensing of targets and the vehicle control hardware.

4.3.1 Range Sensor Recommendations

Range sensors are improving every year. Eaton VORAD already has a new radar unit (EVT-300) that provides not only range and range rate, but also an azimuth angle to the target. This can aid in filtering out targets not in the host's lane of travel. Also, this sensor could be used to extend the experiments to curved portions of road. We recommend acquiring the EVT-300.

A laser-based radar still may be required for the ACC functionality of the system. Time of flight sensors are insensitive to relative velocity and are ideal for this application. The main problem with laser radar is its performance in snow and fog. Snow has proven problematic for laser based range sensors. The Optech range finder was originally designed for mining applications and the company claims that it can sense through dust, which has traditionally been problematic for laser sensors. We need to determine whether this sensor can indeed detect target vehicles in snow. Secondly, no range rate or azimuth data is provided. We can derive range rate from differentiating the range signals, but this was demonstrated to be noisy due to range error (especially at long range) and errors measuring the elapsed time between measurements. Furthermore, the Optech range finder is a point sensor and does not have a beam splitting or a mechanical scanning feature. Optech does sell a mechanical scanner for newer models. We recommend acquiring a laser based range sensor that works in snow and provides range, range rate, azimuth angle to the target. A sensor with these specifications may not exist. In this case, we recommend testing the Optech in snow. If it proves reliable, we can then acquire the latest unit that includes mechanical scanning.

4.3.2 Velocity Controller Hardware and Recommendations

The braking hardware clearly needs improvement. More sensitive control over braking is required so that low levels of velocity can be achieved. Perhaps simply changing the actuator and improving the mechanical advantage of the cable through the addition of pulleys may be sufficient. Otherwise, we may have to try and access the brake controls directly.

The throttle control was adequate for the experiments performed. If there was one thing that could be improved, it would be to have control over the gear shifting. This would allow us to achieve speeds more precisely and avoid jumps in velocity caused by inopportune gear shifting that seemed to trigger oscillations.

Otherwise, the virtual bumper longitudinal algorithm proved stable and effective at controlling the headway distance to the target. Improvements in the sensing and braking hardware will only improve the system. With all the aforementioned hardware deficiencies, the controller achieved the desired headway in a variety of driving situations.

4.3.3 Future Work

The next step in the development of the virtual bumper is to expand its application from only longitudinal control to longitudinal and lateral control. Two options exist for this expansion. The first option is to equip test vehicles with DGPS and radio modems from which targets can continuously broadcast their locations to the host via an RF LAN. This will enable sensing in all directions, providing the host vehicle with complete information regarding the location of all objects to be avoided. The second option is to exploit the capabilities of emerging radar technology (those which provide azimuth information to targets), install enough sensors to provide coverage around the periphery of the host vehicle, and perform experiments with actual hardware. These options will be explored before work is undertaken in this area.

If the LAN approach is taken, after the full system has been tested, we can use the LAN communication system to determine the minimal specifications needed for target sensing. These specifications can then be compared to the most modern range sensors. If needed, we can even have the range sensors custom made to our specifications. Then, a host vehicle can be equipped with these sensors and tested with vehicles not equipped with communication equipment.
If the approach using radar hardware is undertaken, deficiencies in actual radar would be identified. Work with radar vendors could lead to improvements in radar technology, leading to further improvements in the virtual bumper.

References

1. B. Schiller, V. Morellas, and M. Donath. Collision Avoidance for Highway Vehicles Using a Virtual Bumper Controller. *1998 IEEE International Conference on Intelligent Vehicles*, October, 1998.

2. B. Schiller. The Virtual Bumper – A Control Based Collision Avoidance System. M.S. Thesis, University of Minnesota, Dept. of Mechanical Engineering, 1997.

3. Traffic Safety Facts 1997: A Compilation of Motor Vehicle Crash Data from the Fatality Analysis Reporting System and the General Estimates System. National Highway Traffic Safety Administration, November, 1998.

4. S. Shladover. Review of the State of Development of Advanced Vehicle Control Systems (AVCS). *Vehicle System Dynamics*, 1995, pp. 551-595.

5. W. Najm, M. Mironer, J. Koziol Jr., J. Wang, and R. Knipling. Synthesis Report: Examination of Target Vehicular Crashes and Potential ITS Countermeasures. DOT HS 808 263, DOT-VNTSC-NHTSA-95-4, June, 1995.

6. P. Fancher, Z. Bareket, G. Johnson. Predictive analysis of the performance of a headway control system for heavy commercial vehicles. *Proceedings of the 13th IAVSD Symposium on the Dynamics of Vehicles on Roads and Tracks*. Chengdu, China, 1993, pp. 128-141.

7. G. Reichart, K. Naab. Heading Control and Active Cruise Support Driver Assistance Systems for Lateral and Longitudinal Vehicle Guidance. *Proceedings of the 1st World Congress of Transportation Telematics and Intelligent Vehicle Highway Systems*, 4:4126-2133, 1994.

8. K. Naab and G. Reichart. Driver Assistance Systems for Lateral and Longitudinal Vehicle Guidance – Heading Control and Active Cruise Support. *Proceedings of the International Symposium on Advanced Vehicle Control 1994: AVEC '94*, pp. 449-454, October 24, 1994, Tokyo, Japan.

9. H. Satoh and I. Taniguchi. A New Control Method of Adaptive Cruise Control. *Proceedings of the 1st World Congress of Transportation Telematics and Intelligent Vehicle Highway Systems*, pp. 2094-2101, 1994.

10. T. Watanabe, N. Kishimoto, K. Hayafune, K. Yamada, N. Maede. Development of an Intelligent Cruise Control System. *Proceedings of the 2nd World Congress on Intelligent Transport Systems*, November, 1995, Yokohama, Japan.

11. D. Perryman. Tomorrow's Safer Cars. *Newsweek Magazine*, November, 1998, pp. 16.

12. S. Chakraborty and D. Smedley. Adaptive Cruise Control for Heavy Duty Vehicles. *Proceedings of the 1995 annual meeting of ITS America*, Washington, D.C., USA, 1995, pages 145-150.

13. G. Widmann, W. Bauson, and S. Alland. Development of Collision Avoidance Systems at Delphi Automotive Systems. *1998 IEEE International Conference on Intelligent Vehicles*, October 28-30, 1998.

14. H. Fritz. Model-Based Neural Distance Control for Autonomous Road Vehicles. *Proceedings of the 1996 IEEE Intelligent Vehicles Symposium*, pp. 29-34, September 19-20, Tokyo, Japan.

15. H. Kuroda, S. Kuragaki, T. Minowa, K. Nakamura, K. Takano. An Adaptive Cruise Control System Using A Millimeter Wave Radar. *1998 IEEE International Conference on Intelligent Vehicles*, October 28-30, 1998.

16. J. Woll. VORAD Collision Warning Radar. *Proceedings of the IEEE International Radar Conference*, pp. 369-372, May, 1995, Alexandria, Virginia.

17. R. Tribe, K. Prynne and I. Westwood. Intelligent Driver Support. *Proceedings of the 2nd World Congress on Intelligent Transportation Systems*, Vol III, pp. 1187-1192, November 9-11, Yokohama Japan.

18. H. Araki, K. Yamada, Y Hiroshima and T. Ito. Development of Rear-end Collision Avoidance System. *Proceedings of the 1996 IEEE Intelligent Vehicles Symposium*, pp. 224-229, September 19-20, Tokyo, Japan, 1996.

19. T. Butsuen, A. Doi, H. Sasaki. Development of a Collision Avoidance System with Automatic Brake Control. *Proceedings of the 1st World Congress of Transportation Telematics and Intelligent Vehicle Highway Systems*, pp. 2079-2086, 1994.

20. S. Choi and J. Hedrick. Vehicle Longitudinal Control Using an Adaptive Observer for Automated Highway Systems. *Proceedings of the American Control Conference*, pp. 3106-3110, June 1995, Seattle, Washington.

21. U. Ozguner, C. Hatipoglu, K. Unyelioglu, J. Young et al. Radar Based Convoying Using a Frequency Selective Surface Patch. *Proceedings of the 1996 Annual Meeting of ITS America*, pp. 750-755, April 15-18, Houston, Texas.

22. İ. Haskara, C. Hatipoglu, U. Özgüner. Combined Decentralized Longitudinal and Lateral Controller Design for Truck Convoys. *Proceedings of the IEEE Inteligent Transportation Systems Conference*, November 9-12, 1997, Boston, Massachusetts.

23. U. Franke, F. Böttiger, Z. Zomotor, D. Seeberger. Truck Platooning in Mixed Traffic. *Proceedings of the Intelligent Vehicles '95 Symposium*, pp. 1-6, September 25-26, 1995, Detroit, Michigan.

24. P. Borodani, P. Carrea, L. Gortan. Short Headway Control for Trucks on Motorways: the CHAUFFEUR Project. *Proceedings of the 4th World Congress on Intelligent Transport Systems*, October 1997, Berlin, Germany.

25. O. Gehring and H. Fritz. Practical Results of a Longitudinal Control Concept for Truck Platooning with Vehicle to Vehicle Communication. *Proceedings of the IEEE Inteligent Transportation Systems Conference*, November 9-12, 1997, Boston, Massachusetts.

26. M. Hennessey, C. Shankwitz, and M. Donath. Sensor-based Virtual Bumpers for Collision Avoidance: Configuration Issues. *Proceedings of SPIE – The International Society for Optical Engineering*, 2592:48-59, 1995.

27. P. Fancher and Z. Bareket. Controlling Headway Between Successive Vehicles. *Proceedings of the Association for Unmanned Vehicle Systems Annual Meeting*, pp. 149-157, 1994, Detroit, MI.

28. P. Fancher and Z. Bareket. Evaluating Headway Control Using Range Versus Range Rate Relationships. *Vehicle System Dynamics*, pp. 575-596, 1994.

29. T. Gillespie. Fundamentals of Vehicle Dynamics. *Society of Automotive Engineers, Inc.*, 1992.

30. J. Sayer. Intelligent Cruise Control – Issues for Consideration. *Proceedings of the Society of Automotive Engineers Future Transportation Technology Conference*, SAE Technical Paper No. 961667, 1996, Vancouver, Canada.

31. L. Alexander. Differential GPS Based Control of a Heavy Vehicle. M.S. Thesis, University of Minnesota, Dept. of Mechanical Engineering, 1998.

32. R. Kempf. Mvma /nhtsa /sae heavy truck round robin brake test. *SAE Technical Paper Series*, (922484):1-12, 1992.

33. L. Alexander, M. Donath, M. Hennessey, V. Morellas, C. Shankwitz. A Lateral Dynamic Model of a Tractor-Trailer: Experimental Validation. Minnesota Department of Transportation Research Report number 97-18, November, 1996.

34. P. Fancher, R. Ervin, Z. Bareket, G. Johnson, M. Trefalt, J. Tiedecke, and W. Hagleitner. Intelligent Cruise Control: Performance Studies Based Upon an Operating Prototype. *Proceedings of IVHS America 1994 Annual Meeting*, pp. 391-399, 1994, Atlanta, GA.

35. S. Bajikar, A. Gorjestani, P Simpkins, and M. Donath. Evaluation of In-Vehicle GPS-Based Lane Position Sensing for Preventing Road Departure. *Proceedings from the IEEE Conference on Intelligent Transportation Systems*, Boston, Massachusetts, 1997.

36. P. Fancher, J. Sayer and Z. Bareket. A Comparison of Manual Versus Automatic Control of Headway as a Function of Driver Characteristics. *Third Annual World Contress on Intelligent Transport Systems*, 1996, Orlando, FL.